Four Essays in Quantitative Finance

Morten Karlsmark

PhD Thesis

This thesis has been submitted to the PhD School of The Faculty of Science,
University of Copenhagen

Supervisors: Rolf Poulsen, University of Copenhagen
Jesper Fredborg Andreasen, Danske Bank

Submitted: July 31, 2013.

Department of Mathematical Sciences
Faculty of Science
University of Copenhagen

Author: Morten Karlsmark
Love Almqvists väg 8
11253 Stockholm
Sweden
mortenkarlsmark@hotmail.com

Submitted: July 31, 2013.

Assessment Committee: Jens Lund, Nordea Markets
Elisa Nicolato, University of Aarhus
Jesper Lund Pedersen, University of Copenhagen

ISBN: 978-87-7078-984-4

Contents

Preface iii

Abstract v

List of papers vii

Summary ix

Sammenfatning xi

1 Introduction 1

2 SABR Excursions 13

3 A numerical scheme for perpetual claims 41

4 Fast Ninomiya-Victoir calibration of the Double-Mean-Reverting Model 77

5 Utility optimization and feedback-effects 113

Preface

This thesis has been prepared as part of the requirements for the PhD degree at the Department of Mathematical Sciences, Faculty of Science, University of Copenhagen, Denmark. The project was funded by Danske Bank and the Danish Agency for Sciences, Technology and Innovation under the Industrial PhD programme. The work was conducted from August 2010 to July 2013 under supervision of Prof. Rolf Poulsen and Jesper Fredborg Andreasen.

The thesis consist of an introduction and four chapters on different topics. The four chapters are written as academic papers and are thus self-contained. There are minor notational differences among them but none that should cause confusion.

Financial support from Danske Bank and the Danish Agency for Sciences, Technology and Innovation is gratefully acknowledged.

I would like to thank my two advisors Rolf Poulsen and Jesper Fredborg Andreasen for support, ideas and discussions during the work. Furthermore, I wish to thank all members of the Kwant team at Danske Bank and my fellow colleagues at the Department of Mathematical Sciences for discussion on related and unrelated subjects. From September 2011 to December 2011 I was a visiting PhD student at The Department of Mathematics, Baruch College, New York City. I would like to thank Jim Gatheral and the rest of the department for this great experience. I would also like to thank Jim Gatheral and Christian Bayer for co-authoring one of the papers in the thesis.

Last but not least I would like to thank Jenny and my family for their love and support during this work.

Morten Karlsmark
Frederiksberg, July 2013

Abstract

This thesis deals with derivatives, the pricing of them and the effect they can have on financial markets. In the SABR model we develop a new expansion for call prices and a fast arbitrage free pricing scheme that uses a one step implicit finite difference scheme. Next we consider perpetual claims. We suggest a simple finite difference method to price them and prove convergence for the method. We then apply the method to different examples: We calculate the expected exit time of a diffusion from an interval, price perpetual CDOs on households with exploding credit risk, calculate the value of cash when the interest rate can go negative and value perpetual range accruals in the presence of transaction costs. Thirdly we apply modifications of the Ninomiya-Victoir Monte Carlo scheme to the double-mean-reverting (DMR) model (a three factor stochastic volatility model). Thereby we demonstrate on the one hand that fast calibration of the DMR model is practical, and on the other that suitably modified Ninomiya-Victoir schemes are applicable to the simulation of much more complicated time-homogeneous models than may have been thought previously. Last we construct a simple model to capture feedback effects in financial markets. Using the model we can under certain assumptions back out the option position of delta hedgers in a market.

List of papers

This thesis is based on four papers:

- Morten Karlsmark (2012), SABR Excursions, submitted for publication.
- Morten Karlsmark (2013), A numerical scheme for perpetual claims.
- Christian Bayer, Jim Gatheral and Morten Karlsmark (2013), Fast Ninomiya-Victoir calibration of the Double-Mean-Reverting Model.
  Accepted for publication in Quantitative Finance. Available at:
  http://dx.doi.org/10.1080/14697688.2013.818245
- Morten Karlsmark (2013), Utility optimization and feedback-effects.

Summary

This thesis studies different problems in quantitative finance. It consist of four papers that can be read independently. A short description of the papers are given below:

SABR Excursions

Using Markovian projection we construct a local volatility model approximating the SABR model. This local volatility model can be used to build fast approximations to European call prices in the SABR model. We develop a closed form approximation using results from Andersen and Brotherton-Ratcliffe (2005) and an arbitrage-free approximation using one implicit finite difference step. The performance is tested against a numerical solution obtained by finite difference and the approximations from Hagan et al. (2002) and Obloj (2008). The new approximations approximate the numerical solution better and the arbitrage-free approximation eliminates the arbitrage opportunities for low strike options that have become a problem in regimes with very low interest rates.

A numerical scheme for perpetual claims

This paper considers a numerical scheme to price perpetual claims. Contrary to claims with a given maturity date which can be priced solving a PDE, perpetual claims can be priced solving an ODE, often a two boundary problem. We apply a simple finite difference scheme to these two-boundary ODEs. Convergence results are given in a number of cases and we test the method on a variety of different financial examples: We calculate the expected exit time of a diffusion from an interval, price perpetual CDOs on households with exploding credit risk, calculate the value of cash when the interest rate can go negative and value perpetual range accruals in the presence of transaction costs.

Fast Ninomiya-Victoir calibration of the Double-Mean-Reverting Model

We consider the three factor double mean reverting (DMR) model of Gatheral (2008), a model which can be successfully calibrated to both VIX options and SPX options simultaneously. One drawback of this model is that calibration may be slow because no closed form solution for European options exists. In this paper, we apply modified versions of the second order Monte Carlo scheme of Ninomiya and Victoir (2008) and compare these to the Euler-Maruyama scheme with partial truncation of Lord et al. (2010), demonstrating on the one

hand that fast calibration of the DMR model is practical, and on the other that suitably modified Ninomiya-Victoir schemes are applicable to the simulation of much more complicated time-homogeneous models than may have been thought previously.

Utility optimization and feedback-effects
This paper considers a simple model for a market with one asset and two agents. One is optimizing expected utility of future wealth and the other is delta hedging derivatives written on the asset. Using the optimizing behavior of the first agent we can under simple assumptions back out the value of the delta hedgers option position. We develop the ideas first in a continuous time model, where we obtain approximations of the position, then we move on to use the LVI model of Andreasen and Huge (2011). Finally we present numerical examples where the model is applied to option data from different foreign exchange crosses to find the delta hedgers option position.

Bibliography

Andersen, L. and R. Brotherton-Ratcliffe (2005). Extended libor market models with stochastic volatility. *Journal of Computational Finance 9*(1), 1–40.

Andreasen, J. and B. Huge (2011, March). Volatility interpolation. *Risk Magazine 23*(3), 76–79.

Gatheral, J. (2008). Consistent modelling of SPX and VIX options. Bachelier congress, London.

Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002, September). Managing smile risk. *Wilmott Magazine 1*(1), 84–108.

Lord, R., R. Koekkoek, and D. van Dijk (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance 10*(2), 177–194.

Ninomiya, S. and N. Victoir (2008). Weak approximation of stochastic differential equations and application to derivative pricing. *Applied Mathematical Finance 15*, 107–121.

Obloj, J. (2008, May). Fine-tune your smile, correction to Hagan et al. *Wilmott Magazine 35*, 102–104.

Sammenfatning

Denne afhandling studerer forskellige problemer indenfor kvantitativ finansiering. Den består af fire artikler der kan læses uafhængigt. Nedenfor gives en kort introduktion til artiklerne.

SABR Excursions

Ved at bruge Markovian projection konstruerer vi en lokal-volatilitets-model der approksimerer SABR-modellen. Denne lokal-volatilitets-model kan bruges til at finde hurtige approksimationer til Europæiske call-optioner i SABR-modellen. Vi udvikler en approksimation i lukket-form ved at bruge resultater fra Andersen and Brotherton-Ratcliffe (2005) og derefter en arbitragefri approksimation ved at bruge et implicit finite difference skridt. Metoderne testes mod en numerisk løsning, der findes ved en finite difference metode, og approksimationerne i Hagan et al. (2002) and Obloj (2008). De nye approksimationer passer bedre med den numeriske løsning og den arbitragefri approksimation fjerner de arbitrage muligheder der er blevet et problem for optioner med lav strike pris i regimer med meget lav rente.

A numerical scheme for perpetual claims

Denne artikel beskriver en numerisk metode til at prise kontrakter med uendelig løbetid. I modsætning til kontrakter med en given udløbsdato, der prises ved at løse en PDE, kan evige kontrakter prises ved at løse en ODE, ofte et two-boundary problem. Vi løser disse ODEer med en simpel finite difference metode. Konvegens af metoden vises for en række tilfælde og vi tester derefter metoden på forskellige finansielle problemer. Vi udregner det forventede tidspunkt en given diffusions process forlader et interval, priser CDOer uden udløb på husholdninger med eksplosiv kreditrisiko, beregner værdien af kontanter når renten kan blive negativ og prisfastsætter evige range accruels når der er transaktionsomkostninger forbundet med at handle det underlæggende aktiv.

Fast Ninomiya-Victoir calibration of the Double-Mean-Reverting Model

Vi betragter double mean reverting (DMR) modellen fra Gatheral (2008), der succesfuldt kan kalibreres til både VIX og SPX optioner samtidigt. Der er dog en ulempe ved modellen, kalibrering kan være en langsommelig process, da vi ikke har lukkede løsninger til Europæiske optioner. I artiklen anvender vi modificerede udgaver af Ninomiya-Victoir Monte Carlo metoden, se Ninomiya

and Victoir (2008), der er anden ordens konvergent i tid, og holder disse op imod en Euler-Maruyama metode med full truncation fra Lord et al. (2010). Vi demonstrerer herved at DMR modellen kan kalibreres hurtigt og at modificerede udgaver af Ninomiya-Victoir metoden kan anvendes til at simulere langt mere komplicerede tids-homogene modeller end hvad man hidtil har troet.

Utility optimization and feedback-effects
Artiklen opstiller en simpel model for et marked med et aktiv og to agenter. Den ene agent optimerer sin forventede nytte og den anden delta-hedger derivater på aktivet. Ved at bruge den første agents nytteoptimerende adfærd kan vi under simple antagelser regne værdien af den anden agent optionsportefølje ud. Ideerne bliver først udviklet i en kontinuert tids model hvor vi udleder en approksimation til delta-hedgerens optionsposition. Derefter udvikler vi også ideerne i en LVI model, se Andreasen and Huge (2011). Tilsidst præsenteres en række numeriske eksempler, hvor modellen anvendes på optionsdata fra forskellige valutakryds til at finde delta-hedgernes optionsposition.

Bibliography

Andersen, L. and R. Brotherton-Ratcliffe (2005). Extended libor market models with stochastic volatility. *Journal of Computational Finance 9*(1), 1–40.

Andreasen, J. and B. Huge (2011, March). Volatility interpolation. *Risk Magazine 23*(3), 76–79.

Gatheral, J. (2008). Consistent modelling of SPX and VIX options. Bachelier congress, London.

Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002, September). Managing smile risk. *Wilmott Magazine 1*(1), 84–108.

Lord, R., R. Koekkoek, and D. van Dijk (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance 10*(2), 177–194.

Ninomiya, S. and N. Victoir (2008). Weak approximation of stochastic differential equations and application to derivative pricing. *Applied Mathematical Finance 15*, 107–121.

Obloj, J. (2008, May). Fine-tune your smile, correction to Hagan et al. *Wilmott Magazine 35*, 102–104.

Chapter 1

Introduction

1.1 General results in Mathematical finance

In the celebrated paper Black and Scholes (1973) Fischer Black and Myron Scholes used arbitrage arguments to prove a formula for the value of a European call option written on a lognormally distributed asset. This led to an explosion of research in option pricing using arbitrage arguments.

We will briefly present some results on arbitrage free option pricing. They will help us relate the different papers of the thesis to the general theory. They are all taken from Björk (2004).

Let us consider a market with $N+1$ assets $S_t = (S_t^0, \ldots, S_t^N)'$. We assume they follow some diffusion process under the real world probability measure P:

$$\mathrm{d}S_t = \mu(t, S_t)\mathrm{d}t + \sigma(t, S_t)\mathrm{d}W_t^P \tag{1.1.1}$$

A strong mathematical result called "The First Fundamental Theorem", Björk (2004) Theorem 10.14, states that the market is arbitrage free[1] if there exists an equivalent measure Q such that the process $\frac{S_t^i}{S_t^0}$ for all i is a martingale under Q. Another result, "The Second Fundamental Theorem" Björk (2004) Theorem 10.17, states that the model is complete (we can hedge any derivative written on $(S_t)_{t\geq 0}$) if there exists one and only one equivalent martingale measure Q. Using the Girsanov theorem Björk (2004) Theorem 11.3 and the Converse Girsanov Theorem Björk (2004) Theorem 11.6, we know how to change measure from P to any equivalent probability measure. We therefore write the dynamics of S_t under Q by

$$\mathrm{d}S_t = \mu(t, S_t)^Q\mathrm{d}t + \sigma(t, S_t)\mathrm{d}W_t^Q \tag{1.1.2}$$

Let us then expand the market by introducing a derivative written on the assets. Denote the price of this derivative by V_t. We would like the market to

[1] In fact we should say: There exists no free lunch at vanishing risk in the market

be arbitrage free after the introduction of V_t, therefore $\frac{V_t}{S_t^0}$ must be a martingale under one of the equivalent measures Q. If there only exists one equivalent martingale measure then the price is uniquely defined and we can replicate V_t completely by trading in S_t. As $\frac{V_t}{S_t^0}$ is a martingale we have

$$\frac{V_t}{S_t^0} = \mathbb{E}^Q\left(\frac{V_T}{S_T^0}|\mathcal{F}_t\right) \text{ for } T > t$$

if we therefore know the distribution of $\frac{V_T}{S_T^0}$ under Q we can price the claim.

There also exists a PDE representation of the price. If we assume $\frac{V_t}{S_t^0} = f(t, S_t)$ and $\frac{V_T}{S_T^0} = g(S_T)$ then using the Feynman-Kac Theorem, Björk (2004) proposition 5.8, backwards we can write up a PDE for the function $f(t, S_t)$.

$$\begin{aligned} 0 &= f_t(t,s) + \sum_{i=0}^{N} \mu(t,s)_i^Q \frac{\partial f}{\partial s_i}(t,s) + \frac{1}{2}\sum_{i,j=0}^{N} C_{ij}(t,s)\frac{\partial^2 f}{\partial s_i \partial s_j}(t,s) \\ f(T,s) &= g(s) \end{aligned}$$

where $C(t,s) = \sigma(t,s)\sigma(t,s)'$.

We can use any of the two representations to price derivatives, sometimes it is more convenient to use the first sometimes the second.

Another important result is the forward Kolmogorov equation, Björk (2004) proposition 5.12. The density for the process S_t under the Q measure can be found solving the forward PDE

$$\begin{aligned} \frac{\partial p}{\partial t}(t,s) &= -\sum_{i=0}^{N} \frac{\partial}{\partial s_i}\left(\mu_i(t,s)^Q p(t,s)\right) + \frac{1}{2}\sum_{i,j=0}^{N} \frac{\partial^2}{\partial s_i \partial s_j}\left(C_{ij}(t,s)p(t,s)\right) \\ p(0,s) &= \delta(s - s_0) \end{aligned}$$

here we again define $C(t,s) = \sigma(t,s)\sigma(t,s)'$.

If we know the derivative only depends on the density of the underlying at a single time point (the derivative could for example be a call option) then we just need the density of the asset at this point to obtain the price.

1.2 Financial models

Equation (1.1.1) specifies a general financial model in this section we present some different version of it.

The Black-Scholes model is the simplest possible model one can imagine in continuous time finance (neglecting the Bachelier model). We assume

$$\begin{aligned} \mathrm{d}S_t^0 &= rS_t^0 \mathrm{d}t \\ \mathrm{d}S_t^1 &= \mu S_t^1 \mathrm{d}t + \sigma S_t^1 \mathrm{d}W_t^P \end{aligned}$$

the first asset is risk free and is often referred to as the bank account. Using the stated results above we find that under the Q-measure

$$\mathrm{d}S_t^1 = rS_t^1\mathrm{d}t + \sigma S_t^1\mathrm{d}W_t^Q$$

and the derivative price is

$$V_t = \exp(-r(T-t))E^Q(V_T).$$

In this model we can price a number of different derivatives in closed form, these include European call options.

It is well documented that the Black-Scholes model cannot produce prices for derivatives that correspond to the market prices we see. Therefore a large number of other models have been developed, they are normally formulated under the risk neutral measure and we often assume the risk free asset follows the process of S^0 above. Most often these models will be constructed such that the prices for European call options are available in closed form. This is to make calibration of the model easier, we return to this in section 1.4.

We first mention a model which actually lies outside the framework above, this model is the so-called Merton model, see Merton (1976) here

$$dS_t^1 = (r-\lambda m)S_{t-}\mathrm{d}t + \sigma S_{t-}\mathrm{d}W_t^Q + S_{t-}\mathrm{d}J_t^Q$$

where J_t^Q is compound Poisson process with intensity λ and jumps of size $\exp(Y)-1$ for a normal random variable Y. We have defined $m = E\left(\exp(Y)-1\right)$.

Another extension is the CEV model, see for example Schroder (1989)

$$\mathrm{d}S_t^1 = rS_t\mathrm{d}t + \sigma S_t^\beta\mathrm{d}W_t$$

for $\beta \leq 1$.

The three models above will all produce arbitrage free option prices and there exists closed form solutions to the prices of European call options. But even though the Merton model and the CEV model can produce a wider class of option prices than the Black-Scholes model it is not certain that they can fit all option prices in the market. The local volatility model of Dupire (1994) is a non-parametric model that can fit all European call option prices in the market. The model is defined

$$\mathrm{d}S_t = rS_t\mathrm{d}t + \sigma(t,S_t)S_t\mathrm{d}W_t$$

where $\sigma(t,s)$ is a general function chosen such that the model reprices all European call options. The model does not yield closed form formulas for option prices we have to use numerical methods discussed in the next section.

Instead of changing the volatility to a function of S_t we could formulate a model where the volatility is stochastic. The Heston model, see Heston (1993), is specified by

$$\begin{aligned}\mathrm{d}S_t &= rS_t\mathrm{d}t + \sqrt{V_t}S_t\mathrm{d}W_t^1\\ \mathrm{d}V_t &= \kappa(\theta - V_t)\mathrm{d}t + \epsilon\sqrt{V_t}\mathrm{d}W_t^2\end{aligned}$$

and $\mathrm{d}W_t^1 \mathrm{d}W_t^2 = \rho \mathrm{d}t$. Closed form solutions to European call options exists in this model. Extensions of the model can be found in Bates (1996) and in the general framework of Duffie et al. (2000).

We also mention the popular SABR model of Hagan et al. (2002), which is a mixture of a CEV model and a stochastic volatility model

$$\begin{aligned} \mathrm{d}S_t &= rS_t \mathrm{d}t + \alpha_t S_t^\beta \mathrm{d}W_t^1 \\ \mathrm{d}\alpha_t &= \epsilon \alpha_t \mathrm{d}W_t^2 \end{aligned}$$

and $\mathrm{d}W_t^1 \mathrm{d}W_t^2 = \rho \mathrm{d}t$. In this model no closed form solutions for European option prices exists. The reason for its popularity is an expansion formula developed in Hagan et al. (2002) that gives approximative European call prices.

In recent years options written on the volatility of large equity indexes have appeared. Models that can price these options consistently with the options written on the index itself have therefore become popular. We mention Bergomi (2005), Gatheral (2008) and Cont and Kokholm (2013).

1.3 Pricing options using numerical methods

In the Black-Scholes model, the Merton jump model, the CEV model and the Heston model closed form solutions to European call options exists. In the SABR model expansions exist that approximates the option prices. For special local volatility models expansions for European call option values also exists see Hagan and Woodward (1999) and Andersen and Brotherton-Ratcliffe (2005). For the general local volatility case one could use the numerical approximation in Gatheral and Wang (2012).

In other models or for exotic options where closed form solution are not possible we are forced to use numerical methods to price the options.

Let us consider a one dimensional model

$$\mathrm{d}S_t = rS_t \mathrm{d}t + \sigma(t, S_t) S_t \mathrm{d}W_t^Q$$

The option price can either be written as the discounted expectation under the risk neutral measure

$$V_0 = E^Q \left(\exp(-rT) g(S_T) \right)$$

or as the solution to the PDE

$$0 = \frac{\partial f}{\partial t}(t,s) + rs \frac{\partial f}{\partial s}(t,s) + \frac{1}{2}\sigma(t,s)^2 s^2 \frac{\partial^2 f}{\partial s^2}(t,s) - rf(t,s) \tag{1.3.1}$$

$$f(T,s) = g(s) \tag{1.3.2}$$

If we can simulate a stochastic variable with the same distribution as S_T under Q then we can use the law of large numbers to calculate $E^Q \left(\exp(-rT) g(S_T) \right)$

as the average over enough simulated values. The central limit theorem more or less implies that the error will be $O(N^{-1/2})$ where N is the number of simulated values, this error is normally called the integration error. The numerical method presented here is called the Monte Carlo method. In the Black-Scholes model S_T will be lognormally distributed, therefore by simulating enough lognormally distributed variables we can calculate European option prices. In most cases though the distribution of S_T will not be known, we only know the SDE that generates the random variable. In order to simulate S_T we discretize the SDE, for example using an Euler-Maruyama scheme

$$S_{t_{i+1}} = rS_{t_i}\Delta t + \sigma(t_i, S_{t_i})S_{t_i}\sqrt{\Delta t}Z_{i+1}$$

where Z_i are iid standard normally distributed. One can show that this discretization leads to an error of order $O(\Delta t)$, this error is normally called the discretization error. So we have two error types, integration error and discretization error.

The integration error can normally be reduced using Quasi Monte Carlo methods for example using Sobol numbers, see Joe and Kou (2008).

There also exists discretization methods that lead to lower discretization error. We could for example use the Ninomiya-Victoir scheme, see Ninomiya and Victoir (2008). If we consider a one dimensional time homogeneous diffusion

$$\mathrm{d}S_t = \frac{1}{2}\sigma(S_t)\sigma'(S_t)\mathrm{d}t + \sigma(S_t)\mathrm{d}W_t$$

The corresponding Backward PDE for our model becomes

$$\frac{\partial u}{\partial \tau}(\tau, x) = \frac{1}{2}\sigma^2 u(\tau, x)$$

where $\tau = T - t$ and we have defined $\sigma u(\tau, x) = \sigma(x)\frac{\partial u}{\partial x}(\tau, x)$. By differentiating in τ we obtain

$$\frac{\partial^2 u}{\partial \tau^2}(\tau, x) = \frac{1}{4}\sigma^4 u(\tau, x)$$

We can therefore write

$$u(\Delta\tau, x) = u(0, x) + \frac{1}{2}\Delta\tau\sigma^2 u(0, x) + \frac{1}{8}\Delta\tau^2\sigma^4 u(0, x) + O(\Delta\tau^3)$$

Define $\exp(V)x = g(1)$ where

$$\frac{dg}{dt} = V(g), \quad g(0) = x.$$

Set $f(x) = u(0, x)$ and let us calculate $E\left(f\left(\exp\left(\sqrt{\Delta\tau}Z\sigma\right)x\right)\right)$ where Z is standard normally distributed. After a long tedious calculation one gets

$$E\left(f\left(\exp\left(\sqrt{\Delta\tau}Z\sigma\right)x\right)\right) = \left(1 + \frac{1}{2}\Delta\tau\sigma^2 + \frac{1}{8}\Delta\tau^2\sigma^4\right)f(x) + O(\Delta\tau^3)$$

which implies

$$E\left(f\left(\exp\left(\sqrt{\Delta\tau}Z\sigma\right)x\right)\right) = u(\Delta\tau, x) + O(\Delta\tau^3).$$

So the scheme gives third order convergence for one time step. We normally consider a fixed maturity so we need $O\left(\frac{1}{\Delta\tau}\right)$ time steps therefore we have a second order scheme.

This scheme will be used in chapter 4 to speed up option pricing in the model of Gatheral (2008).

Instead of using a Monte Carlo method we can solve the pricing PDE (1.3.1) by Finite Differences. Consider the PDE

$$\begin{aligned} 0 &= \frac{\partial f}{\partial t}(t,s) + \frac{1}{2}\sigma(t,s)^2\frac{\partial^2 f}{\partial s^2}(t,s) \\ f(T,s) &= g(s) \end{aligned}$$

we want to approximate the function $f(t,s)$ with a function $v(t,s)$ defined on an equidistant discrete grid $\{t_0,\ldots,t_M\}\times\{s_0,\ldots s_N\}$. We set the initial condition:

$$v(t_M, s_j) = g(s_j) \text{ for } s_j \in \{s_0,\ldots s_N\}$$

and approximates the PDE in the internal points $s_j \in \{s_1,\ldots s_{N-1}\}$

$$\begin{aligned} 0 = \frac{v(t_i,s_j)-v(t_{i-1},s_j)}{\Delta t} &+ (1-\theta)\frac{1}{2}\sigma(t_i,s_j)^2\frac{v(t_i,s_{j+1})-2v(t_i,s_j)+v(t_i,s_{j-1})}{\Delta x^2} \\ &+ \theta\frac{1}{2}\sigma(t_{i-1},s_j)^2\frac{v(t_{i-1},s_{j+1})-2v(t_{i-1},s_j)+v(t_{i-1},s_{j-1})}{\Delta x^2} \end{aligned}$$

here θ determines how we approximate the spatial derivative. On the boundary we could use $v(t_{i-1},s_0) = v(t_i,s_0)$ and $v(t_{i-1},s_N) = v(t_i,s_N)$ as an approximation. We collect the equations for all the grid points and obtain a linear equation system

$$(I - \theta\Delta t A_{t_{i-1}})v(t_{i-1}) = (I + (1-\theta)\Delta t A_{t_i})v(t_i)$$

this system can be solved very efficiently for $v(t_{i-1})$ using the tridag() algorithm of Press et al. (2002). When $\theta = 0.5$ the scheme will be second order convergent both in space and in time. $\theta = 1$ gives a scheme which in itself can be seen as discrete time discrete state space model with non-negative transition densities. This model will always produce arbitrage free prices.

Chapter 3 considers the use of finite difference methods to perpetual claims and chapter 2 uses finite difference methods both to obtain prices in the SABR model but also to obtain arbitrage free approximations to option prices using a one step finite difference approximation with $\theta = 1$.

The discrete time discrete state space model when $\theta = 1$ is used in chapter 5 to calibrate a local volatility model and to solve for a utility maximizing investors optimal portfolio.

1.4 The use of financial models in practice

Having chosen a model to price and risk manage our derivatives we need to specify the parameters of the model. This is done by calibrating the model to liquid derivatives for which market prices exists. The calibration procedure is normally formulated as quadratic minimization problem, let $V_i(\Omega)$ be the model price of derivative i that depends on the parameters Ω, and let $\tilde{V}_i$ be the market quote. We then wish to solve the problem

$$\min_{\Omega} \sum_{i=1}^{N} \left(f_i(V_i(\Omega)) - f_i(\tilde{V}_i) \right)^2 \tag{1.4.1}$$

where f_i is a function such that the problem is specified in the right quantity (perhaps we want to solve the problem in dollars or perhaps in implied Black volatilities).

The quadratic problem is solved using a numerical solver such as the Levenberg-Marquart algorithm, see Press et al. (2002). The solver evaluates the model prices of the derivatives a large number of times and searches for the minimum. In order for the calibration to be fast we therefore need the calculation of derivative prices to be fast. Both chapter 2 and 4 focus on fast calculation of option prices to make a fast calibration possible.

When we have obtained the parameters Ω^* that minimizes the quadratic problem (1.4.1), we can use the model to price non-standard derivatives, these will be priced consistently with the liquid derivatives in the market. But just as important, having specified a model for the underlying we know how the prices of the derivatives moves with the prices of the underlying and other derivatives. Therefore we know how to hedge derivatives using the underlying and possibly other more liquid derivatives.

Imagine a trader having sold a derivative with value V_t. His model specifies the value as a function of the underlying (and possibly other processes) $V_t = f(t, S_t)$. To hedge himself against fluctuations in the price V_t the trader will delta hedge the derivative ie he will buy $\frac{\partial f}{\partial s}(t, S_t)$ of the underlying.

In all models considered up until now, we assume the agents in the market can trade any amount of the underlying without affecting the price of it. This will of course never be true. If we knew the trading patterns of the agents in the market we would know when the drift and volatility of the asset price would be high and low. We do not know the trading patterns, but every option trader knows his own option book and therefore how he will delta hedge it in the future. We could imagine that this information is being processed by the option market and therefore that the option prices expresses among other things the position of the delta hedgers in the market. Chapter 5 will build a model to extract the delta hedgers option position from option prices.

1.5 The next chapters

Chapter 2 deals with the SABR model. The expansion formula for call prices from Hagan et al. (2002) is not arbitrage free, low and high strike options are priced badly. Using Markovian projection and results from Doust (2012) we develop a local volatility model that approximately produces the same call prices as the SABR model. We can then use expansion results from Andersen and Brotherton-Ratcliffe (2005) to price call options fast. The new expansion will produce fever arbitrage opportunities than the one from Hagan et al. (2002) but it will not be arbitrage free. A fast arbitrage free approximation is then developed, we use the local volatility model obtained and apply a one step implicit finite difference scheme to it together with a correction term, this will produce arbitrage free prices. In the SABR model the local volatility function is a CEV function. We can in fact use our approximation for other functions. We therefore also apply our results to a SABR model with a hyperbolic volatility function, see Jackel and Kahl (2008), and obtain an expansion with very few arbitrage opportunities.

Chapter 3 deals with perpetual claims. If the underlying model is time-homogeneous then the price of a perpetual claim is the solution to a second order ODE (often a two boundary problem). We apply a simple finite difference scheme to solve the ODEs and prove convergence result for linear and nonlinear ODEs. In the end we consider different examples of financial interest.

Chapter 4 considers the second order Monte Carlo scheme of Ninomiya and Victoir (2008) and the application of this to the double-mean-reverting model of Gatheral (2008). The Ninomiya-Victoir scheme relies on fast solution of multidimensional ODEs specified by the model we like to simulate. The double-mean-reverting model gives rise to ODEs that cannot be solved in closed form. We extend the Ninomiya-Victoir scheme such that the ODEs of the new scheme can be solved in closed form. By this we illustrate how the scheme can be modified to yield closed form simulation for other more complicated models. Next we apply the scheme to the calibration of the double-mean-reverting model and compare the performance to an Euler-Maruyama scheme with partial truncation of Lord et al. (2010). We conclude that the Ninomiya-Victoir scheme can lead to an increase in performance and that we can calibrate the double-mean-reverting model fast. We also show numerically that the scheme exhibits second order convergence in time.

Chapter 5 builds a simple feedback-effects model in a two person economy. The first agent is a utility optimizing investor and the second is an option trader delta hedging his option position. We assume the asset has one price when the first agent is alone in the market and another when both agents are present. The second price is assumed to be a smooth function of the first. Using these assumptions and the dynamics of the price under the equivalent martingale measure we can approximate the dynamics of the price under the real world measure. Having found that we can calculate the first agents demand for the

asset, which will give us the delta hedgers demand and thereby his option position. We construct a continuous time model and also a discrete time discrete state space model. Numerical examples are provided, here we apply the discrete model to option data from different foreign exchange crosses.

Bibliography

Andersen, L. and R. Brotherton-Ratcliffe (2005). Extended libor market models with stochastic volatility. *Journal of Computational Finance 9*(1), 1–40.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate process implicit in deutsche mark option. *The Review of Financial Studies 9*(1), 69–107.

Bergomi, L. (October 2005). Smile dynamics II. *Risk Magazine 18*(10), 67–73.

Björk, T. (2004). *Arbitrage theory in continuous time* (Second ed.). Oxford University Press.

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy 81*(3), 637–654.

Cont, R. and T. Kokholm (2013). A consistent pricing model for index options and volatility derivatives. *Mathematical Finance 23*(2), 248–274.

Doust, P. (2012). No-arbitrage SABR. *Journal of Computational Finance 15*(3), 3–31.

Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica 68*(6), 1343–1376.

Dupire, B. (1994, January). Pricing with a smile. *Risk 7*, 18–20.

Gatheral, J. (2008). Consistent modelling of SPX and VIX options. Bachelier congress, London.

Gatheral, J. and T.-H. Wang (2012). The heat-kernel most-likely-path approximation. *International Journal of Theoretical and Applied Finance 15*(1), 1250001.

Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002, September). Managing smile risk. *Wilmott Magazine 1*(1), 84–108.

Hagan, P. S. and D. E. Woodward (1999). Equivalent black volatilities. *Applied Mathematical Finance 6*(3), 147–157.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility and applications to bond and currency options. *The Review of Financial Studies 6*, 327–343.

Jackel, P. and C. Kahl (2008, March). Hyp hyp hooray. *Wilmott Magazine 34*, 70–81. New version is available at: http://www.christiankahl.com/publications/HypHypHooray.pdf.

Joe, S. and F. Y. Kou (2008). Constructing sobol sequences with better two-dimensional projections. *SIAM J. Sci. Comput 30*, 2635–2654.

Lord, R., R. Koekkoek, and D. van Dijk (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance 10*(2), 177–194.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics 3*, 125–144.

Ninomiya, S. and N. Victoir (2008). Weak approximation of stochastic differential equations and application to derivative pricing. *Applied Mathematical Finance 15*, 107–121.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (2002). *Numerical recipes in C++* (Second ed.). Cambridge University Press.

Schroder, M. (1989). Computing the constant elasticity of variance option pricing formula. *The Journal of Finance 44*(1), 211–219.

Chapter 2

SABR Excursions

Morten Karlsmark

Abstract

Using Markovian projection we construct a local volatility model approximating the SABR model. This local volatility model can be used to build fast approximations to European call prices in the SABR model. We develop a closed form approximation using results from Andersen and Brotherton-Ratcliffe (2005) and an arbitrage-free approximation using one implicit finite difference step. The performance is tested against a numerical solution obtained by finite difference and the approximations from Hagan et al. (2002) and Obloj (2008). The new approximations approximate the numerical solution better and the arbitrage-free approximation eliminates the arbitrage opportunities for low strike options that have become a problem in regimes with very low interest rates.

2.1 Introduction

The SABR model is used widely by financial institutions to price and risk-manage interest rate derivatives such as European swaptions and Constant Maturity Swaps. It is a stochastic volatility model where the forward F_t (for a specific expiry) follows

$$\begin{aligned} \mathrm{d}F_t &= \alpha_t F_t^{\beta} \mathrm{d}W_t^1 \\ \mathrm{d}\alpha_t &= s\alpha_t \mathrm{d}W_t^2 \end{aligned} \tag{2.1.1}$$

here $\mathrm{d}\left\langle W^1, W^2 \right\rangle_t = \rho dt$. Note that α_t is lognormally distributed, and the process for the forward F_t is of the CEV type. The last fact implies that the forward has a non-zero probability of hitting zero. If it hits zero it will stay there by arbitrage arguments, see Andersen and Andreasen (2000).

No closed form solution exists for European claims but we can solve the problem numerically using finite difference, see the appendix in section 2.9, or

Monte Carlo. These approaches are though to slow to be used in practice as traders need to risk manage products with hundreds of different underlying instruments in different currencies. Therefore practitioners have been using an approximation from Hagan et al. (2002) to value European options. This approximation, though, introduces arbitrage opportunities in the market, low and high strike options are priced incorrectly. This has led researchers to suggest a number of other approximations. Most of these approximations lead to arbitrage possibilities, others are very computationally involved. This paper suggests a new, computationally fast approximation that generates arbitrage-free prices. The method uses Markovian projection to obtain a local volatility model that approximates the SABR model. From this model we develop one closed form approximation using results from Andersen and Brotherton-Ratcliffe (2005) and an arbitrage-free approximation using a one step implicit finite difference scheme. The results are general and they can be applied to other local volatility functions than the CEV function of the normal SABR model.

The paper is organized as follows. Section 2.2 considers the Hagan approximation and provides numerical examples in order to demonstrate the inherent problems with the approximation. Section 2.3 looks at new approximation methods that have appeared in the literature. In section 2.4 we develop a local volatility model that approximates the SABR model, this is done using Markovian projection. Section 2.5 approximates call prices in the local volatility model, these are not arbitrage-free. Section 2.6 develops a one step implicit finite difference scheme to solve for prices in the local volatility model. This scheme generates arbitrage free prices. The results in section 2.4 and 2.5 are very general, in section 2.7 we apply them to a model where the CEV function is replaced by a function with bounded derivative at 0. Section 2.8 concludes.

2.2 The Hagan Approximation

The approximation found in Hagan et al. (2002) has made the SABR model popular compared to other stochastic volatility models because it is easy to use and program. We here give a short review of the approximation and presents some numerical examples.

Consider the process $((F_t - K)^+)_{t\geq 0}$ and let us find dynamics for it. Define $f(x) = (x-K)^+$, then we have $f'(x) = 1_{\{x>K\}}(x)$ and $f''(x) = \delta(x-K)$. Where $\delta(\cdot)$ is a Dirac delta function. Using Tanaka's formula we get

$$(F_T - K)^+ = (F_0 - K)^+ + \int_0^T 1_{\{F_t>K\}}(F_t)\alpha_t F_t^\beta \mathrm{d}W_t + \frac{1}{2}\int_0^T \delta(F_t - K)\alpha_t^2 F_t^{2\beta}\mathrm{d}t.$$

Taking mean and using Tornelli we obtain

$$\begin{aligned}
\mathbb{E}\left((F_T-K)^+\right) &= (F_0-K)^+ + \frac{1}{2}\int_0^T \mathbb{E}\left(\delta(F_t-K)\alpha_t^2 F_t^{2\beta}\right)\mathrm{d}t \\
&= (F_0-K)^+ + \frac{1}{2}\int_0^T \int\int \delta(F_t-K)\alpha_t^2 F^{2\beta} p(t,F_t,\alpha_t)\mathrm{d}F_t\mathrm{d}\alpha_t\mathrm{d}t \\
&= (F_0-K)^+ + \frac{1}{2}\int_0^T \int \alpha_t^2 K^{2\beta} p(t,K,\alpha_t)\mathrm{d}\alpha_t\mathrm{d}t \\
&= (F_0-K)^+ + \frac{K^{2\beta}}{2}\int_0^T \int \alpha_t^2 p(t,K,\alpha_t)\mathrm{d}\alpha_t\mathrm{d}t
\end{aligned}$$

here $p(t,F_t,\alpha_t)$ is the two dimensional probability density for (F_t,α_t). The derivation here can be proven rigorously, see Carr and Jarrow (1990).

Hagan et al. (2002) uses perturbation techniques in order to approximate $\int \alpha_t^2 p(t,K,\alpha_t)\mathrm{d}\alpha_t$. Let us denote this approximation $A(t,\beta,s,F_0,\alpha_0,K)$. Then they approximate $\int_0^T A(t,\beta,s,F_0,\alpha_0,K)\mathrm{d}t$ with an integral $\int_{f(\beta,s,F_0,\alpha_0,K)}^{\infty} g(y)\mathrm{d}y$ where $g(y)$ is independent of β,s,F_0,α_0 and K. Setting $f(0,0,F_0,\sigma,K)=f(\beta,s,F_0,\alpha_0,K)$ and solving for σ they obtain an equivalent normal volatility for any combination of β,s,F_0,α_0 and K. This normal volatility is then transformed into a Black volatility.

But the Hagan approximation has some problems. Let us compare the Black volatilities found using our finite difference approach in section 2.9 with the Black volatilities from the Hagan approximation. This is done in Figure 2.1 where we picture Black volatilities for 10Yx10Y Swaptions using parameters from the 15th of December 2010. We assume that the swap-rate follows (2.1.1) under the annuity measure. As seen the Hagan approximation clearly overvalues all options. This could be acceptable. If we use the Hagan approximation to calibrate our model to the market and then price using the Hagan approximation, the model prices will correspond to the market prices in the neighborhood of the market quotes. The true problem is that the prices obtained by the Hagan approximation deep in and out of the money are not arbitrage-free. This can be seen if we differentiate the swaption prices twice in the strike direction to obtain the Arrow-Debreu prices, see Figure 2.2. The Arrow-Debreu prices make up the probability density of the 10Yx10Y forward swap-rate under the annuity measure. For the finite difference solution the density shoots up when the strike approaches zero, this reflects the fact that the forward has a positive probability of ending up at 0. The Hagan expansion on the other hand gives a negative density for low strikes, this leads to arbitrage opportunities:

Just consider a normal butterfly. Buy one call option at strike x_1, one at strike x_2 and sell two options at strike $\frac{x_1+x_2}{2}$ (we assume $x_1<x_2$). The price of this strategy approximates the second derivative of the call price in the strike direction, modulo a multiplication of a positive constant. If x_1 and

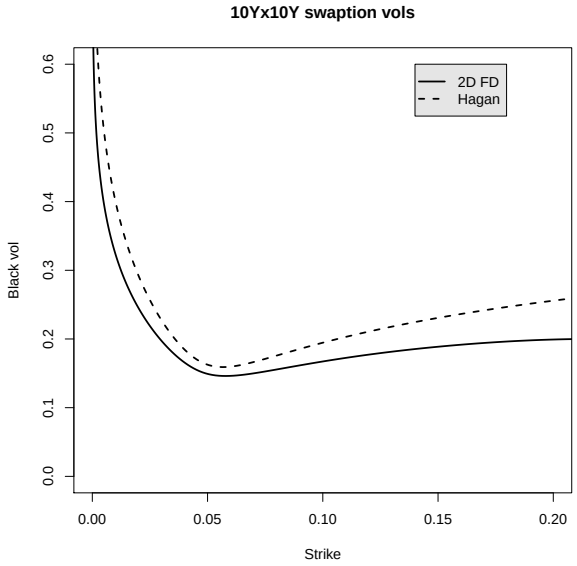

Figure 2.1: 10Yx10Y swaption volatilities for the SABR model found using a 2D FD scheme and the Hagan approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

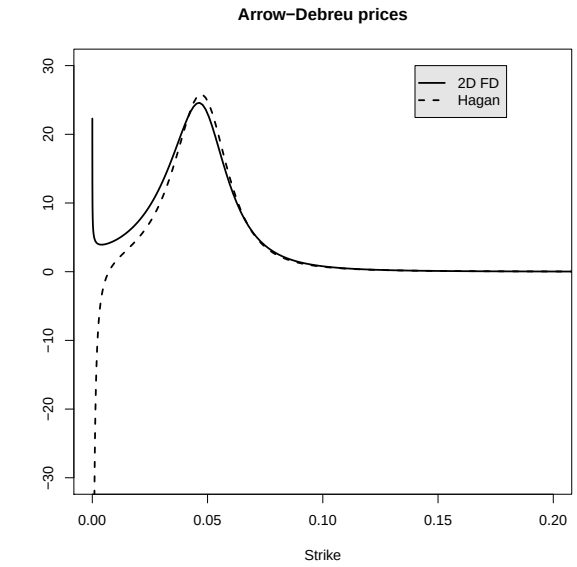

Figure 2.2: Arrow-Debreu prices for the 10Yx10Y forward swap rate in SABR model. Found using a 2D FD scheme and the Hagan approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

x_2 are low strikes we will obtain a negative price for such a strategy. But the payoff will always be a non-negative amount. Hence we have constructed an arbitrage opportunity.

To avoid these arbitrage opportunities other approximations have appeared in the literature.

2.3 Other approximations

Obloj (2008) compares a result from Berestycki et al. (2004) with the result in Hagan et al. (2002). He notes a slight difference between the formulas and then shows that the formula from Berestycki et al. (2004) is better at calculating the prices for low strike calls. But it is worth noting that Hagan et al. (2002) actually develops two different expressions for the Black volatility. A general formula for any form of local volatility function and a specific formula for the CEV case. The first general formula can also be used in the SABR case, but it gives worse prices than the next specific formula. It is actually the first general formula that Obloj (2008) considers. The Obloj (2008) formula also gives rise to arbitrage for low strikes, and it produces a very steep smile as the formula from Hagan et al. (2002).

As noted Hagan et al. (2002) approximates the time integral $\int_0^T A(t, \beta, s, F_0, \alpha_0, K)\mathrm{d}t$. Benhamou and Croissant (2007) note that this integral can be solved explicitly using the complex error function. Therefore their result only relies on the accuracy of the approximation $A(t, \beta, s, F_0, \alpha_0, K)$. This method is a huge improvement compared to Hagan et al. (2002) and Obloj (2008). But we cannot guarantee in general that the prices will be arbitrage-free, since we still rely on the approximation of $\int \alpha_t^2 p(t, K, \alpha_t)\mathrm{d}\alpha_t$. Also we need to evaluate the complex error function, this will make the method time consuming.

Another path is chosen in Doust (2012). He derives approximations for $\int \alpha_t^n p(t, K, \alpha_t)\mathrm{d}\alpha_t$ for any n. This is done using the same kind of perturbation techniques as Hagan et al. (2002). $\int p(t, K, \alpha_t)\mathrm{d}\alpha_t$ is the marginal density for F_t as a function of K. The approximation to the marginal density is positive so if we use this as the density we get arbitrage-free call prices. But the approximated density does not integrate to one, and the mean is not equal to the forward price. This means that we need to scale and shift the approximated density to make the model work. The obtained prices may be arbitrage-free but it comes at the cost of computation time.

The methods above are derived using expansions. In contrast Kennedy et al. (2012) takes a probabilistic approach to the problem. They develop a displaced diffusion SABR model approximating the CEV SABR model. In this model they obtain closed form expressions for the mean and variance of the forward given the stochastic volatility. Using these they approximate the conditional distribution of the forward given the stochastic volatility with a

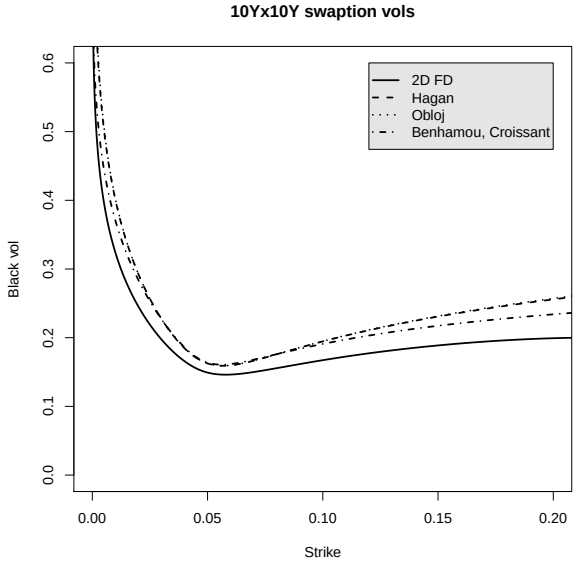

Figure 2.3: 10Yx10Y swaption volatilities for the SABR model found using a 2D FD scheme, the Hagan approximation, the Obloj approximation and the Benhamou and Croissant approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

normal distribution or a NIG distribution. As the stochastic volatility is log-normally distributed, approximative prices can be found by a two dimensional integral and prices will be arbitrage-free. For the normal approximation further simplifications leads to a method where only one integral has to be evaluated. This method produces very nice results, but the integral can diverge for large vol of vol and maturity. The method will probably also be slower than the methods above, since we have to compute an integral numerically.

Hagan et al. (2002), Obloj (2008) and Benhamou and Croissant (2007) smiles can be seen in Figure 2.3. As seen the Hagan et al. (2002) and Obloj (2008) smiles look very alike while the Benhamou and Croissant (2007) smile lies closer to the true solution.

2.4 The new method

The idea is to develop a local volatility model that prices call options as the SABR model. We therefore want a model of the form

$$\mathrm{d}F_t = \sigma(F_t, t)\mathrm{d}W_t \tag{2.4.1}$$

(note we use the normal model form here) and we need to choose $\sigma(F_t, t)$ such that the call prices from the local volatility model agree with the prices from

the SABR model. Derman and Kani (1998) shows that this can be achieved if we choose

$$\sigma(K,t) = K^{\beta}\sqrt{\mathbb{E}(\alpha_t^2|F_t = K)}, \tag{2.4.2}$$

this is also known as Markovian projection, see Piterbarg (2007). Therefore we need to compute

$$\mathbb{E}(\alpha_t^2|F_t = K) = \frac{\int \alpha_t^2 p(t,K,\alpha_t)\mathrm{d}\alpha_t}{\int p(t,K,\alpha_t)\mathrm{d}\alpha_t}. \tag{2.4.3}$$

The numerator is the function Hagan et al. (2002) approximates and as mentioned Doust (2012) approximates functions of the form $\int \alpha_t^n p(t,K,\alpha_t)\mathrm{d}\alpha_t$ for any n. We can therefore use the results from Doust (2012) to derive the conditional second moment.

2.4.1 Doust's result

Consider the model

$$\begin{aligned} \mathrm{d}F_t &= \alpha_t C(F_t)\mathrm{d}W_t^1 \\ \mathrm{d}\alpha_t &= s\alpha_t \mathrm{d}W_t^2 \end{aligned} \tag{2.4.4}$$

where $\mathrm{d}\left\langle W^1, W^2\right\rangle_t = \rho \mathrm{d}t$ and $C(\cdot)$ is some function.

Doust (2012) shows that

$$\begin{aligned} \int \alpha_t^n p(t,K,\alpha_t)\mathrm{d}\alpha_t \approx \frac{\alpha_0^{n-1}}{C(K)}\sqrt{\frac{C(F_0)}{C(K)}}J(y(F_0,K))^{n-3/2} \\ \cdot \exp\left(\int_0^{y(F_0,K)} \frac{1/2\rho s\alpha_0 \frac{B'(\alpha_0\xi,K)}{B(\alpha_0\xi,K)}\xi}{J(\xi)^2}\mathrm{d}\xi\right)\frac{\exp(-\frac{x(y(F_0,K))^2}{2t} + \kappa t)}{\sqrt{2\pi t}} \end{aligned}$$

where

$$\begin{aligned} J(y) &= \sqrt{1 - 2\rho s y + s^2 y^2} \\ y(F_0,K) &= \frac{1}{\alpha_0}\int_K^{F_0} \frac{1}{C(f)}\mathrm{d}f \\ x(y) &= \int_0^y \frac{1}{J(\xi)}\mathrm{d}\xi \\ \kappa &= \frac{1}{8}(4n^2 - 8n + 2 - (2n-3)(2n-1)\rho^2)s^2 \\ &\quad + \frac{2n-1}{4}\rho s\alpha_0 \frac{B'(\alpha_0 y_0,K)}{B(\alpha_0 y_0,K)} + \alpha_0^2\left(\frac{1}{4}\frac{B''(\alpha_0 y_0,K)}{B(\alpha_0 y_0,K)} - \frac{3}{8}\left(\frac{B'(\alpha_0 y_0,K)}{B(\alpha_0 y_0,K)}\right)^2\right) \\ B'(x,K) &= \frac{\partial B}{\partial x}(x,K). \end{aligned}$$

$B(x,K)$ is defined by the equality $B(\alpha_0 y(F_0,K),K) = C(F_0)$ and y_0 is a constant we have to choose. The result is true for any choice of y_0, we will therefore for each value of F_0 and K choose $y_0 = y(F_0,K)$. This means that y_0 is not constant in F_0 and K but it won't matter. The interested reader is referred to appendix A in Doust (2012).

2.4.2 The conditional second moment

Dividing the approximation for $\int \alpha_t^2 p(t,K,\alpha_t)\mathrm{d}\alpha_t$ with the approximation for $\int p(t,K,\alpha_t)\mathrm{d}\alpha_t$ we obtain

$$\mathbb{E}(\alpha_t^2|F_t = K) \approx \alpha_0^2 J(y(F_0,K))^2 \exp\left(\rho s\alpha_0 \frac{B'(\alpha_0 y_0,K)}{B(\alpha_0 y_0,K)} t\right) \tag{2.4.5}$$

and for the SABR model where $C(F) = F^\beta$ we have

$$\begin{aligned} J(y) &= \sqrt{1-2\rho s y + s^2 y^2} \\ y(F_0,K) &= \begin{cases} \frac{1}{\alpha_0}\log(\frac{F_0}{K}) & \beta = 1 \\ \frac{1}{\alpha_0(1-\beta)}(F_0^{1-\beta} - K^{1-\beta}) & \beta < 1 \end{cases} \\ B(y,K) &= \begin{cases} K\exp(y) & \beta = 1 \\ (y(1-\beta)+K^{1-\beta})^{\frac{\beta}{1-\beta}} & \beta < 1. \end{cases} \end{aligned}$$

Note that the right hand side of (2.4.5) is always positive, so we can use it as a local variance.

2.4.3 Call prices

European call prices can be found solving Dupire's forward equation

$$\frac{\partial C}{\partial T}(T,K) = \frac{1}{2}\sigma(T,K)^2 \frac{\partial^2 C}{\partial K^2} \tag{2.4.6}$$

where C is the call price and

$$\sigma(T,K)^2 = K^{2\beta}\alpha_0^2 J(y(F_0,K))^2 \exp\left(\rho s\alpha_0 \frac{B'(\alpha_0 y_0,K)}{B(\alpha_0 y_0,K)} T\right). \tag{2.4.7}$$

If we use an implicit finite difference scheme the prices will be arbitrage-free see Andreasen and Huge (2011b). Prices from a 1D Finite Difference scheme can be seen in Figure 2.4.

2.5 An approximative method

Solving the local volatility model using a finite difference scheme is, like solving the 2-dimensional problem, too time consuming to be used in practice. Instead

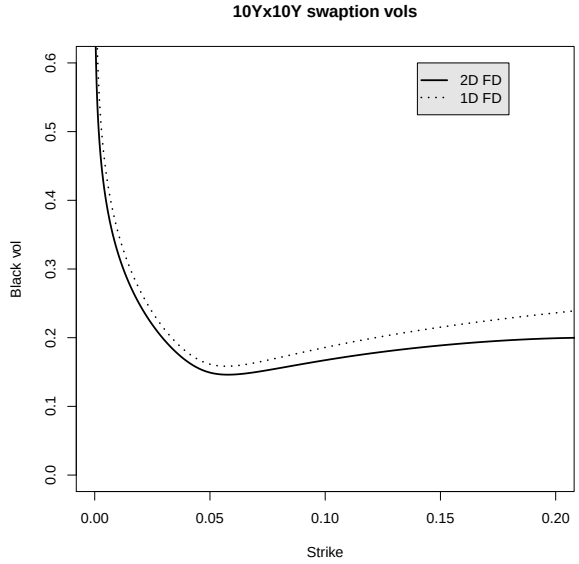

Figure 2.4: 10Yx10Y swaption volatilities for the SABR model found using a 2D FD scheme and the 1D FD scheme. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

we now develop an approximation to call prices using a method from Andersen and Brotherton-Ratcliffe (2005).

The price V of an European option in our local volatility model can be found solving the 1D backward PDE:

$$\frac{\partial V}{\partial t}(t,F) + \frac{1}{2}\sigma(t,F)^2\frac{\partial^2 V}{\partial F^2}(t,F) = 0. \tag{2.5.1}$$

We can write

$$\sigma(t,F)^2 = \psi(F_0,F)^2\lambda(t) \tag{2.5.2}$$

where

$$\psi(F_0,F) = F^\beta \alpha_0 J(y(F_0,F)) \quad \text{and} \quad \lambda(t) = \exp\left(\rho s\alpha_0 \frac{B'(\alpha_0 y_0, K)}{B(\alpha_0 y_0, K)} t\right). \tag{2.5.3}$$

Therefore we can make a deterministic time shift: Define

$$u(t) = \int_0^t \lambda(s)ds \tag{2.5.4}$$

and we have the new PDE

$$\frac{\partial V}{\partial u}(u,F) + \frac{1}{2}\psi(F_0,F)^2\frac{\partial^2 V}{\partial F^2}(u,F) = 0 \tag{2.5.5}$$

because $u'(t) = \lambda(t)$. Andersen and Brotherton-Ratcliffe (2005) show that we can approximate the price of a strike K maturity $u(T)$ call option with the Black formula using the volatility

$$\sigma_u = \Omega_0(F_0, K) + \Omega_1(F_0, K)u(T)$$

where

$$\begin{aligned}
\Omega_0(F_0, K) &= \frac{\log(F_0/K)}{\int_K^{F_0} \psi(F_0, x)^{-1} dx} \\
\Omega_1(F_0, K) &= -\frac{\Omega_0(F_0, K)}{\left(\int_K^{F_0} \psi(F_0, x)^{-1} dx\right)^2} \log\left(\Omega_0(F_0, K)\sqrt{\frac{KF_0}{\psi(F_0, K)\psi(F_0, F_0)}}\right).
\end{aligned}$$

This is when $u(t)$ is the measure of time, inverting to t as our measure we obtain

$$\sigma(F_0, K)_{Black} = \frac{\Omega_0(F_0, K)u(T)^{1/2} + \Omega_1(F_0, K)u(T)^{3/2}}{\sqrt{T}}. \tag{2.5.6}$$

As seen the denominators of $\Omega_0(F_0, K)$ and $\Omega_1(F_0, K)$ goes to zero as $K \to F_0$. We could therefore experience numerical instability when K comes close to F_0. Using the rule of l'Hôpital we find that

$$\begin{aligned}
\lim_{K \to F_0} \Omega_0(F_0, K) &= \frac{\psi(F_0, F_0)}{F_0} \\
\lim_{K \to F_0} \Omega_1(F_0, K) &= \frac{1}{24}\frac{\psi(F_0, F_0)}{F_0} \\
&\quad \cdot \left(\frac{\psi(F_0, F_0)^2}{F_0^2} - \frac{\partial \psi}{\partial K}(F_0, F_0)^2 + 2\psi(F_0, F_0)\frac{\partial^2 \psi}{\partial K^2}(F_0, F_0)\right).
\end{aligned}$$

We can therefore use these limits when K goes to F_0. In order to guarantee numerical stability around the forward we simply interpolate the smile linearly between $F_0 - 0.0001$ and $F_0 + 0.0001$.

Then we calculate $\int_K^{F_0} \psi(F_0,x)^{-1}dx$.

$$\begin{aligned}
\int_K^{F_0} \psi(F_0,x)^{-1}dx &= \int_K^{F_0} \left(x^\beta \alpha_0 J\left(y(F_0,x)\right)\right)^{-1} dx \\
&= \frac{1}{\alpha_0}\int_K^{F_0} x^{-\beta}(1-2\rho s y(F_0,x)+s^2 y(F_0,x)^2)^{-1/2}dx \\
&= -\int_{y(F_0,K)}^{y(F_0,F_0)} (1-2\rho s y+s^2y^2)^{-1/2}dy \\
&= -\left[\frac{1}{s}\log\left(\left|2s\sqrt{1-2\rho s y+s^2y^2}+2s^2y-2\rho s\right|\right)\right]_{y(F_0,K)}^{y(F_0,F_0)} \\
&= \frac{1}{s}\log\left(\left|2s\sqrt{1-2\rho s y(F_0,K)+s^2y(F_0,K)^2}\right.\right. \\
&\qquad\qquad \left.\left.+2s^2y(F_0,K)-2\rho s\right|\right) - \frac{1}{s}\log\left(|2s-2\rho s|\right) \\
&= \frac{1}{s}\log\left(\left|\frac{\sqrt{1-2\rho s y(F_0,K)+s^2y(F_0,K)^2}+sy(F_0,K)-\rho}{1-\rho}\right|\right).
\end{aligned} \tag{2.5.7}$$

where we have used

$$y(F_0,K) = \begin{cases} \frac{1}{\alpha_0}\log(\frac{F_0}{K}) & \beta = 1 \\ \frac{1}{\alpha_0(1-\beta)}(F_0^{1-\beta}-K^{1-\beta}) & \beta < 1 \end{cases}$$

and $y(F_0,F_0)=0$.

(2.5.6) is a simple approximation to the Black volatility for European call/put option prices in the SABR model. We can compare it to the results from the 1D finite difference scheme considered above, this is done in Figure 2.5. As seen the approximation is almost as good as solving for option prices in the local volatility model using finite difference, although some differences exists for low strikes. We can also compare it to the approximation from Hagan et al. (2002) this is done in Figure 2.6. The new approximation seems to approximate the call prices better than the Hagan approximation. But like the Hagan approximation it creates arbitrage opportunities in the market, see Figure 2.7.

2.6 Fast arbitrage-free prices

Let us reconcile. We have approximated the two-dimensional SABR model with a one-dimensional local volatility model using Markovian projection. We can find European option prices in this local volatility model by finite difference methods, this will give arbitrage-free prices. We have also developed an

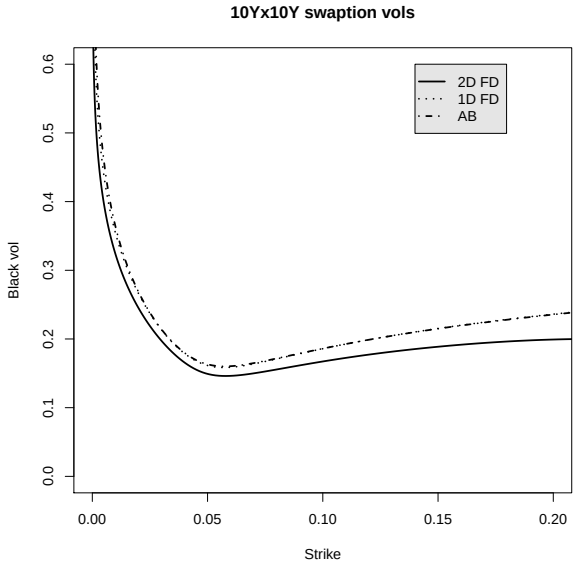

Figure 2.5: 10Yx10Y swaption volatilities for the SABR model found using a 2D FD scheme, the 1D FD scheme and the Andersen Brotherton-Ratcliffe approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

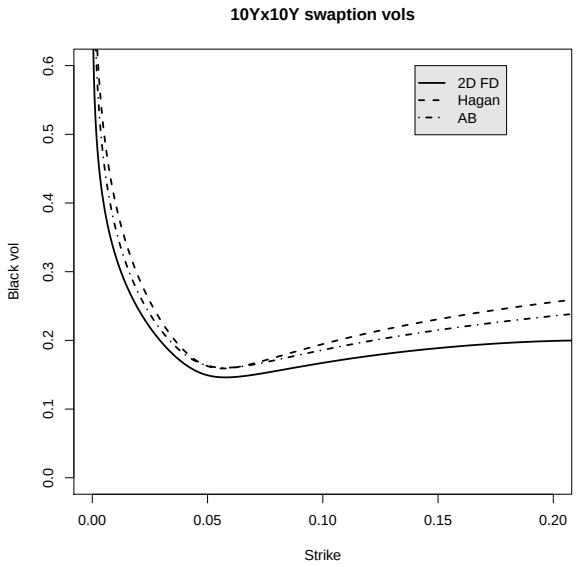

Figure 2.6: 10Yx10Y swaption volatilities for the SABR model found using a 2D FD scheme, the Hagan approximation and the Andersen Brotherton-Ratcliffe approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

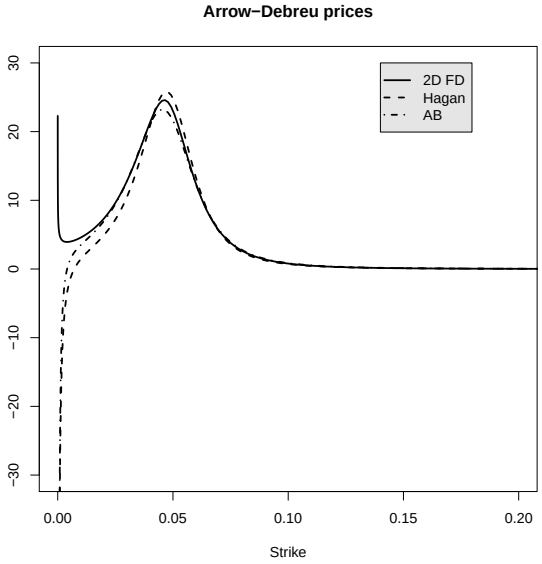

Figure 2.7: Arrow-Debreu prices for the 10Yx10Y forward swap rate in the SABR model. Found using a 2D FD scheme, the Hagan approximation and the Andersen Brotherton-Ratcliffe approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

approximation to call prices in the local volatility model using results from Andersen and Brotherton-Ratcliffe (2005).

Now we are looking for a fast way to calculate arbitrage-free prices, and the idea is to use an implicit finite difference scheme with only one time step. This will ensure arbitrage-free prices, see Andreasen and Huge (2011b). In order to limit the number of grid points in the spatial dimension we will only solve the PDE in a small interval around ATM where options trade. As boundary conditions we will use our Andersen Brotherton-Ratcliffe approximation.

We need to choose the grid width for our finite difference scheme. The width of the grid is chosen as a number of standard deviations in a log-normal distribution. We choose to set the log-normal volatility equal to the ATM implied volatility from our Andersen Brotherton-Ratcliffe approximation. The lower and upper boundary are floored/capped and we construct the grid in log space. Then we transform all points back to normal space. In the following graphs we have $S_{min} = 0.000205$ and $S_{max} = 0.203$ as boundaries of our finite difference grid. In between these points we solve the PDE corresponding to our local volatility model. This procedure will therefore generate arbitrage-free prices in the interval $[S_{min}; S_{max}]$, and on the borders the prices will equal the Andersen Brotherton-Ratcliffe approximation. Outside the interval we use the Andersen Brotherton-Ratcliffe approximation as our price. Implied volatilities calculated using this method are found in Figure 2.8. As seen our volatility

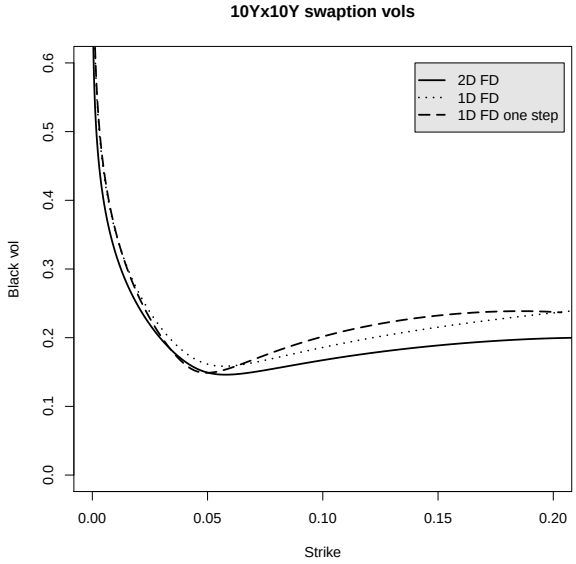

Figure 2.8: 10Yx10Y swaption volatilities for the SABR model found using a 2D FD scheme, 1D FD scheme and our 1D FD one step scheme. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

smile is v-shaped compared to the true solution. The v-shape can be found in almost any model where one uses an implicit finite difference scheme with one time step. It exists because the implicit finite difference scheme with one time step has not got the time steps to smooth out the probability mass. In fact if we calculate the transition density for one implicit finite difference step, we will see it approximates a Laplace distribution when the volatility is constant e.g. in a normal model, see Andreasen and Huge (2011a). In Figure 2.9 we have pictured implied volatilities for a normal model. One smile computed using one time step and another using 100 time steps. The smile is v-shaped when using one time step.

In order to eliminate the v-shape from the smile of the one step scheme we will correct the local volatilities to reflect that we only use one time step. This is the same approach as in Andreasen and Huge (2013). As noted above almost all models yield a v-shaped volatility smile when one uses an implicit finite difference scheme with one time step. Therefore we can calculate a correction to the local volatilities in a model with closed form vanilla prices. This correction can then be used in our local volatility model. We choose to calculate the correction in the normal model (which is the SABR model where $\beta = 0$ and $s = 0$) this is most convenient.

An implicit finite difference scheme with one time step can generate correct

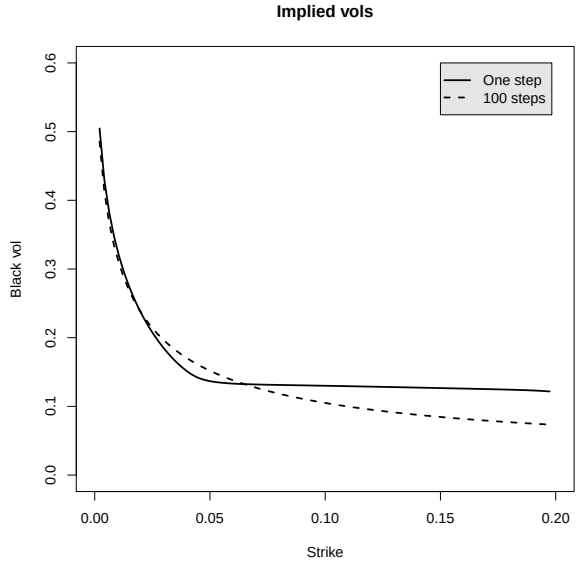

Figure 2.9: Implied Black volatilities for a normal model found using an implicit FD scheme with one time step and an implicit FD scheme with 100 time steps. Parameters are $\sigma = 0.0071$, $F_0 = 0.0439$, $T = 10$.

call prices if the local volatility function satisfy

$$\sigma(T,K)^2 = 2\frac{C(T,K) - C(0,K)}{\frac{\partial^2 C}{\partial K^2}(T,K)_{FD}T} \tag{2.6.1}$$

for any point K. Here $C(T,K)$ is the correct price of a call with expiry T and strike K and $\frac{\partial^2 C}{\partial K^2}(T,K)_{FD}$ is the finite difference for the second derivative on our discrete grid.

The price in a normal model for an expiry T, strike K call is

$$C(T,K) = (F_0 - K)\Phi\left(\frac{F_0 - K}{\sigma\sqrt{T}}\right) + \sigma\sqrt{T}\varphi\left(\frac{F_0 - K}{\sigma\sqrt{T}}\right) \tag{2.6.2}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function and $\varphi(\cdot)$ is the standard normal density function. We also have

$$C(0,K) = (F_0 - K)^+. \tag{2.6.3}$$

We choose to approximate $\frac{\partial^2 C}{\partial K^2}(T,K)_{FD}$ with $\frac{\partial^2 C}{\partial K^2}(T,K)$ which is the density of the risk neutral distribution, i.e. a normal density. So in order for the one step scheme to generate approximately correct call prices in a normal model

we need to have

$$\begin{aligned}\sigma(T,K)^2 &= 2\frac{(F_0-K)\Phi\left(\frac{F_0-K}{\sigma\sqrt{T}}\right)+\sigma\sqrt{T}\varphi\left(\frac{F_0-K}{\sigma\sqrt{T}}\right)-(F_0-K)^+}{\frac{1}{\sigma\sqrt{T}}\varphi\left(\frac{F_0-K}{\sigma\sqrt{T}}\right)T}\\ &= 2\sigma^2\left(1+\frac{(F_0-K)\Phi\left(\frac{F_0-K)}{\sigma\sqrt{T}}\right)-(F_0-K)^+}{\sigma\sqrt{T}\varphi\left(\frac{F_0-K}{\sigma\sqrt{T}}\right)}\right)\\ &= 2\sigma^2\left(1-x\frac{\Phi(-x)}{\varphi(x)}\right)\end{aligned}$$

where $x=\frac{|F_0-K|}{\sigma\sqrt{T}}$. The volatility correction $\frac{\sigma(T,K)}{\sigma}$ is then given by

$$\sqrt{2\left(1-x\frac{\Phi(-x)}{\varphi(x)}\right)}. \tag{2.6.4}$$

$\frac{\Phi(-x)}{\varphi(x)}$ can be approximated by a fifth degree polynomial (see equation 26.2.17 in Abramowitz and Stegun (1972)).

To correct our local volatility model we multiply the previously found local volatility $K^\beta\sqrt{\mathbb{E}(\alpha_T^2|F_t=K)}$ with (2.6.4) at every grid node. But we need to choose the σ in the expression for x. In order for the correction to work we need to choose approximately the correct normal volatility σ for each grid point.

We use the Andersen Brotherton-Ratcliffe approximation to calculate Black volatilities for each grid point, and then we normalize with $\sqrt{F_0K}$

$$\sigma_{NormalK}=\sqrt{F_0K}\sigma_{BlackK}. \tag{2.6.5}$$

This scheme yields the implied volatilities found in Figure 2.10. As seen we get a nice volatility smile, and the method generates arbitrage-free prices in the interval $[S_{min},S_{max}]$, see Figure 2.11. The method is also computationally fast since we use a 1D Finite Difference scheme with one time step.

2.7 Other volatility functions: The Hyperbolic volatility function

In the SABR model the process for the forward is of CEV type. So the forward has a non-zero probability of hitting zero. This could be one of the reasons for the problems of expansion methods. To eliminate the problem we could use a $C(\cdot)$ function with a bounded slope at zero. This function must also have the same characteristics as the CEV function around ATM since we like the

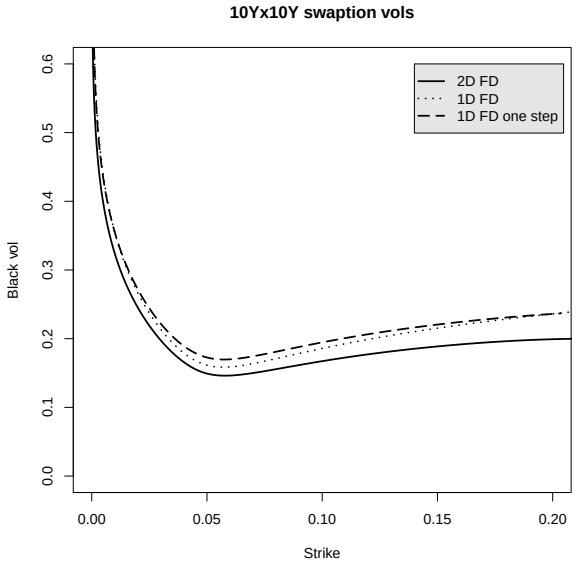

Figure 2.10: 10Yx10Y swaption volatilities for the SABR model found using a 2D FD scheme, our 1D FD scheme and our 1D FD one step scheme with volatility correction. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

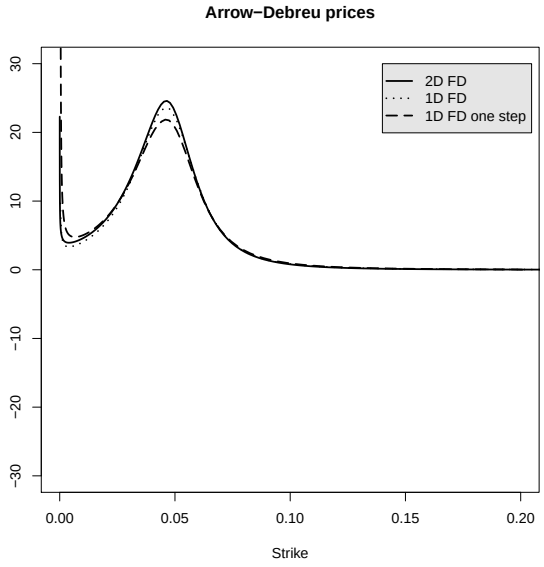

Figure 2.11: Arrow-Debreu prices for the 10Yx10Y forward swap rate in the SABR model. Found using a 2D FD scheme, our 1D FD scheme and our 1D FD one step scheme with volatility correction. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

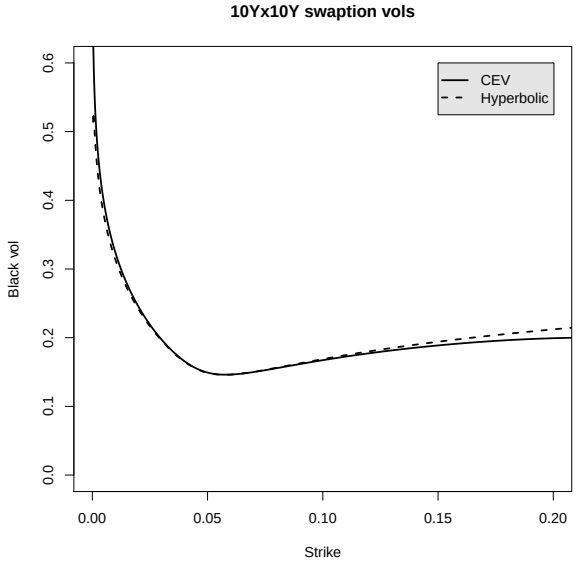

Figure 2.12: 10Yx10Y swaption volatilities for the model with CEV function and Hyperbolic function. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

smiles to look the same. A function that fulfills this is the hyperbolic volatility function

$$C(F) = F_0^{\beta} \frac{(1-\beta+\beta^2)\frac{F}{F_0} + (\beta-1)\left(\sqrt{\left(\frac{F}{F_0}\right)^2 + \beta^2\left(1-\frac{F}{F_0}\right)^2} - \beta\right)}{\beta} \tag{2.7.1}$$

see Jackel and Kahl (2008). It is easily seen that it has the same value, first and second derivative as the CEV function ATM. Therefore we expect it to yield similar option prices as the CEV function around ATM. Figure 2.12 shows implied volatilities found using a finite difference scheme on the SABR model with a CEV function and a hyperbolic function. As seen the volatilities match each other around ATM. For high strikes we see a steeper smile in the hyperbolic model while the opposite is true for low strikes.

Using the general result (2.4.5) we can obtain a local volatility model approximating the SABR model with the hyperbolic volatility function. Again we need to calculate $\int_K^{F_0} \psi(F_0, x)^{-1} \mathrm{d}x$.

$$\begin{aligned}\int_K^{F_0} \psi(F_0,x)^{-1}\mathrm{d}x &= \int_K^{F_0} (C(x)\alpha_0 J(y(F_0,x)))^{-1}\mathrm{d}x \\ &= -\int_{y(F_0,K)}^{y(F_0,F_0)} J(y)^{-1}\mathrm{d}y.\end{aligned}$$

as $y(F_0, K) = \frac{1}{\alpha_0}\int_K^{F_0} \frac{1}{C(f)}\mathrm{d}f$. We can use (2.5.7) and we only need to calculate $\int_K^{F_0} \frac{1}{C(f)}\mathrm{d}f$.

$\int_K^{F_0} \frac{1}{C(f)}\mathrm{d}f$ can be found using Mathematica and we have

$$\int_K^{F_0} \frac{1}{C(F)}\mathrm{d}F = F_0^{1-\beta}\left(\frac{\beta}{2}\log\left(\frac{\beta - \beta\frac{K}{F_0} + \xi(\frac{K}{F_0}))}{(\frac{K}{F_0})^2}\right)\right.$$
$$- \frac{(-1+\beta-\beta^2+\beta^3)}{\beta\sqrt{1+\beta^2}}\log\left(\frac{1+\sqrt{1+\beta^2}}{\beta^2(-1+\frac{K}{F_0})+\frac{K}{F_0}+\sqrt{1+\beta^2}\xi(\frac{K}{F_0})}\right)$$
$$\left.+\frac{1-\beta+\frac{1}{2}\beta^2}{\beta}\log\left(\frac{\beta^2}{\beta^3(-1+\frac{K}{F_0})+\beta^2(2-2\frac{K}{F_0}+\xi(\frac{K}{F_0}))+2(1-\beta)(-\frac{K}{F_0}+\xi(\frac{K}{F_0}))}\right)\right)$$

where $\xi(x) = \sqrt{\beta^2(-1+x)^2 + x^2}$.

Now we can use a finite difference scheme to solve for prices or approximate the prices using the Andersen Brotherton-Ratcliffe result. If our hypothesis is correct the Andersen Brotherton-Ratcliffe approximation should yield fewer arbitrage opportunities when we use a hyperbolic volatility function than when we use a CEV function.

As Hagan et al. (2002) derives a formula for a general $C(F)$ we can in fact also calculate the option prices using their formula. We can therefore compare the different approaches. Figure 2.13 shows prices in a model with hyperbolic volatility function computed using the different approximations. Again the Andersen Brotherton-Ratcliffe approximation yields a better approximation than the Hagan approximation.

In Figure 2.14 we picture Arrow-Debreu prices found using the Andersen Brotherton-Ratcliffe approximation on the model with hyperbolic volatility function and on the model with CEV function. As seen the model with hyperbolic volatility function yields a much nicer density. It gets negative at a much lower strike than the density found in the model with CEV function. The use of the hyperbolic volatility function can therefore reduce the problems of call price expansions in the SABR model. We reach the same conclusion using the Hagan approximation, see Figure 2.15.

But note that the general Hagan approximation is unstable. For high strike options the price increases with strike. This is in fact also the case for the Andersen Brotherton-Ratcliffe approximation but this only happens for extremely high strikes.

2.8 Conclusion

Using Markovian projection and a result from Doust (2012) we developed a local volatility model approximating the SABR model. This local volatility model was then used to construct two approximations to European call option prices in the SABR model. The first approximation was developed using

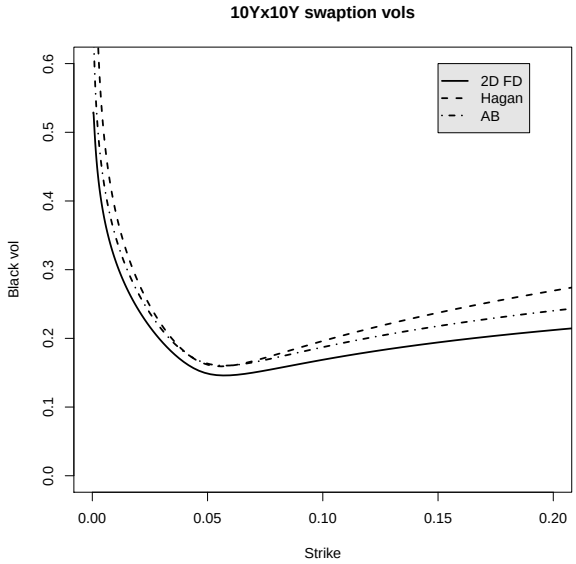

Figure 2.13: 10Yx10Y swaption volatilities for the model with Hyperbolic volatility function. We have used a 2D FD solver, the Hagan approximation and the Andersen Brotherton-Ratcliffe approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

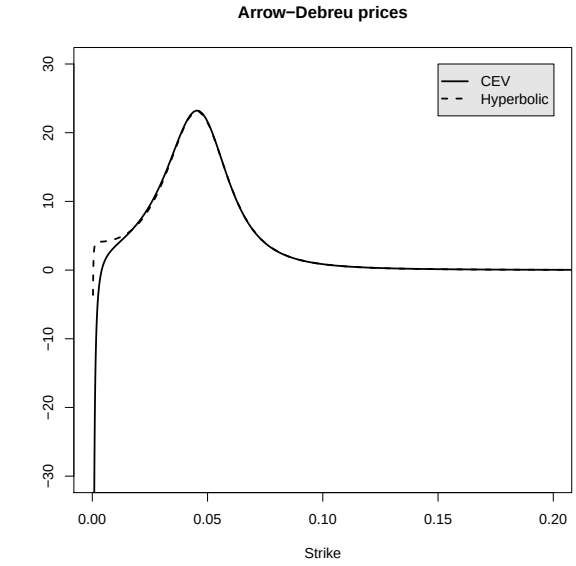

Figure 2.14: Arrow-Debreu prices for the 10Yx10Y forward swap rate in the model with a CEV volatility function and a Hyperbolic volatility function. Found using the Andersen Brotherton-Ratcliffe approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

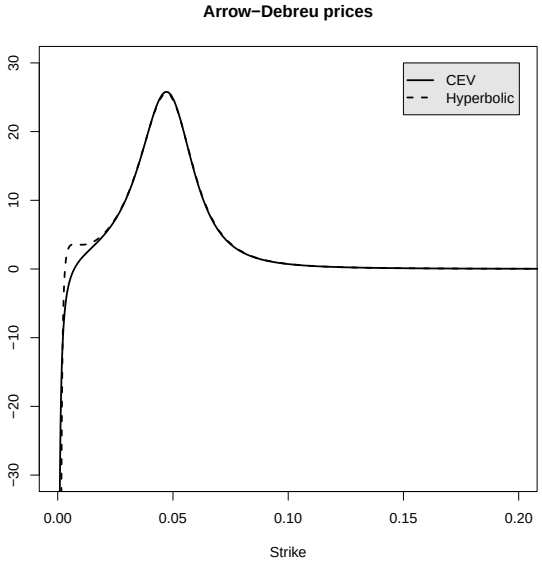

Figure 2.15: Arrow-Debreu prices for the 10Yx10Y forward swap rate in the model with a CEV volatility function and a Hyperbolic volatility function. Found using the Hagan approximation. Parameters are $\beta = 0.7$, $F_0 = 0.0439$, $\alpha_0 = 0.0630$, $s = 0.421$, $\rho = -0.363$.

results from Andersen and Brotherton-Ratcliffe (2005). As shown in numerical examples it was better than the Hagan et al. (2002) approximation but it did not eliminate arbitrage opportunities for low strike options. The second approximation was a one step implicit finite difference solution of the local volatility model. To make it work we had to develop a correction to the local volatility function. The resulting scheme was fast and generated arbitrage-free prices. The Markovian projection result is general and works for any local volatility function we choose for our SABR model. We therefore developed an approximation to a SABR model where the local volatility function was the hyperbolic local volatility function. In this model the forward cannot hit zero. The Andersen Brotherton-Ratcliffe approximation for this model generated prices with fewer arbitrage opportunities than the Andersen Brotherton-Ratcliffe approximation for the normal SABR model. It can therefore be an idea to use a local volatility function with a bounded derivative at zero if one wants easy implementable approximations with fewer arbitrage opportunities.

Acknowledgement

The author would like to thank Jesper Andreasen and Brian Huge for discussions, comments and ideas during this work.

2.9 Appendix: FD solution of the SABR model

2.9.1 The forward FD scheme

The benchmark prices for our SABR model will be calculated using a two dimensional finite difference scheme. Here we explain how. For an introduction to Finite Difference schemes in finance see Tavella and Randall (2000).

Consider a diffusion process $(X_t)_{t\geq 0}$. Let $p(t,x,T,X)$ be the transition density from (t,x) to (T,X). Here t,T is time and x,X is the value of our spatial variables. Also let $V(t,x)$ be the price of a European option at time t when the spatial variable is equal to x. When pricing European options using PDE's we can choose two different approaches. Either we price using the backward Kolmogorov PDE

$$\frac{\partial V}{\partial t}(t,x) + A_x V(t,x) = 0 \tag{2.9.1}$$

or we find a density using the forward Kolmogorov PDE

$$-\frac{\partial p}{\partial T}(t,x,T,X) + A_X p(t,x,T,X) = 0 \tag{2.9.2}$$

and then calculate prices integrating the payoff function with respect to the found density. Here A_x is the infinitesimal generator in the spatial coordinates x and A_X is the adjoint operator in the coordinates X.

Solving the forward PDE will therefore give us all call prices at once. In contrast the Backward PDE needs to be solved for each option individually. If we want prices for many different options the fast method will be to solve the forward equation.

But solving the forward equation directly is not easy. Especially choosing the right boundary conditions for our grid is hard. Here we will use a rewriting of a discretized backward equation and obtain a discretized forward equation.

We follow Andreasen (2009). Consider the normal backward equation:

$$0 = \frac{\partial V}{\partial t} + A_x V$$

where again A_x is the infinitesimal generator. Normally we discretize the spatial dimension and approximate A_x with a matrix $\bar{A}_x$, so we obtain a system

$$0 = \frac{\partial v}{\partial t} + \bar{A}_x v$$

where v is a vector. Now multiply by another vector p and integrate from 0 to T

$$0 = \int_0^T p(t)' \frac{\partial v}{\partial t}(t)dt + \int_0^T p(t)' \bar{A}_x v(t)dt$$

using partial integration on the first part yields

$$\begin{aligned} 0 &= p(T)'v(T) - p(0)'v(0) + \int_0^T p(t)'\bar{A}_x v(t) - \frac{\partial p}{\partial t}(t)'v(t)dt \\ &= p(T)'v(T) - p(0)'v(0) + \int_0^T v(t)'(\bar{A_x}'p(t) - \frac{\partial p}{\partial t}(t))dt. \end{aligned}$$

So if $p(t)$ solves the forward equation

$$\frac{\partial p}{\partial t}(t) = \bar{A_x}'p(t) \tag{2.9.3}$$

and if

$$p(0) = 1_{\{f=F_0,\alpha=\alpha_0\}}(f,\alpha) \tag{2.9.4}$$

then $p(T)$ will be the probabilities for the discrete backward problem. We then have a discrete forward equation for p:

$$\frac{\partial p}{\partial t}(t) = \bar{A_x}'p(t)$$

with initial condition $p(0) = 1_{\{f=F_0,\alpha=\alpha_0\}}(f,\alpha)$. This can be solved with a finite difference scheme. The rewriting from backward to forward equation can be done with any number of spatial dimensions. Note that we only need to specify the $\bar{A}_x$ matrix for the backward problem, which normally is easily done.

We need to find prices in the SABR model which has two spatial dimensions. Therefore we will use a Craig-Sneyd solver, see Craig and Sneyd (1988): For a parabolic forward equation in 2 dimensions

$$\frac{\partial V}{\partial t}(t,x,y) = D_x V + D_y V + D_{xy} V$$

we choose θ_x, θ_y, θ_{xy} and approximate:

$$\begin{aligned} (1-\theta_x\Delta t\bar{D}_x)U &= (1+(1-\theta_x)\Delta t\bar{D}_x + \Delta t\bar{D}_y + \Delta t\bar{D}_{xy})V(t) \\ (1-\theta_y\Delta t\bar{D}_y)\tilde{V}(t+\Delta t) &= U - \theta_y\Delta t\bar{D}_y V(t) \\ (1-\theta_x\Delta t\bar{D}_x)U &= (1+(1-\theta_x)\Delta t\bar{D}_x + \Delta t\bar{D}_y + (1-\theta_{xy})\Delta t\bar{D}_{xy})V(t) \\ &\quad + \theta_{xy}\Delta t\bar{D}_{xy}\tilde{V}(t+\Delta t) \\ (1-\theta_y\Delta t\bar{D}_y)V(t+\Delta t) &= U - \theta_y\Delta t\bar{D}_y V(t) \end{aligned}$$

here $\bar{D}_x, \bar{D}_y$ and $\bar{D}_{xy}$ are matrices approximating the spatial derivatives. In our case we define the matrices $\bar{A}_x$, $\bar{A}_y$ and $\bar{A}_{xy}$ and then set $\bar{D}_x = \bar{A}'_x$, $\bar{D}_y = \bar{A}'_y$ and $\bar{D}_{xy} = \bar{A}'_{xy}$. We will choose $\theta_x = \theta_y = \theta_{xy} = \frac{1}{2}$ which gives second order convergence in time.

2.9.2 Call prices

The scheme will return a matrix p of probabilities in the F and α direction. We can easily calculate the marginal F probabilities by summing over the α direction:

$$P_{F_T=x} = \sum_{i=1}^{m} p_{x,i}$$

here m is the number of points in the α direction.

Let the grid have n points in the F direction, we denote the value of F at each grid point by $(F_i)_{1\leq i\leq n}$. We are interested in call options with strikes in $(F_i)_{1\leq i\leq n}$. To calculate these we simply define

$$Q_j = \sum_{i=j+1}^{n} P_i \quad \text{and} \quad R_j = \sum_{i=j+1}^{n} F_i P_i. \tag{2.9.5}$$

A strike F_j call then has the price

$$C_j = \sum_{i=j+1}^{n} (F_i - F_j)P_i = R_j - F_j Q_j.$$

Starting from $j = n-1$ we can calculate all call prices in $O(n)$ time.

2.9.3 The grid

In our implementation we choose to log transform the α-process such that it will be normally distributed. The width of the grid is chosen to -10 to 10 standard deviations.

We have three points that need to lie on the grid. The lower bound l, the upper bound u and α_0. We know that l will be point 1 and u will be point m on the grid. α_0 will normally not lie on the grid if we do equal spacing. But we can easily find the grid point in an equal spacing that lie closest to α_0, let us call this point r. We now have three pairs of points $(1, l)$, (r, α_0), (m, u). We then do a cubic spline interpolation of these three pairs. To calculate our grid points we insert numbers from 1 to m into this spline function. This ensures that α_0 is among the grid points.

To choose the width of the grid in the F direction we need to approximate the standard deviation of F. We can calculate the standard deviation for the process with α fixed when $\beta = 0$, 0.5 and 1. The standard deviation for a general β is then approximated by a second order polynomial hitting the three points. The width of the grid is chosen to be the interval -10 to 10 standard deviations floored by 0.

To distribute the points on the grid we transform the lower boundary, upper boundary and F_0 using the function $f(x) = x^{1-\beta}$. Then we lay out the

grid points using the cubic spline approach given above. All points are then transformed back to the normal space using the function $g(x) = x^{\frac{1}{1-\beta}}$. The transformation is done in order to distribute the grid points in a nice way such that they vary with the probability distribution.

2.9.4 Boundary conditions

Note that we only need to specify the $\bar{A}_x$ matrix, which is used for the backward scheme. Therefore we just need to choose boundary conditions for the backward scheme. We choose to make the function linear at the boundary ie we set the second order derivatives equal to zero on the boundary. This fits well with the absorption condition when $F_t = 0$.

Bibliography

Abramowitz, M. and I. A. Stegun (1972). *Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables*. Washington, D.C.: National Bureau of Standards. Electronic version found at: http://people.math.sfu.ca/~cbm/aands/.

Andersen, L. and J. Andreasen (2000). Volatility skews and extensions of the Libor market model. *Applied Mathematical Finance 7*(1), 1–32.

Andersen, L. and R. Brotherton-Ratcliffe (2005). Extended libor market models with stochastic volatility. *Journal of Computational Finance 9*(1), 1–40.

Andreasen, J. (2009). Planck-fokker boundary conditions. Working paper.

Andreasen, J. and B. Huge (2011a, July). Random grids. *Risk Magazine 24*(7), 66–71.

Andreasen, J. and B. Huge (2011b, March). Volatility interpolation. *Risk Magazine 24*(3), 76–79.

Andreasen, J. and B. Huge (2013, January). Expanded forward volatility. *Risk Magazine 26*(1), 101–107.

Benhamou, E. and O. Croissant (2007). Local time for the SABR model, connection with the "complex" Black Scholes and application to CMS and Spread Options. Working paper available at SSRN: http://ssrn.com/abstract=1064461.

Berestycki, H., J. Busca, and I. Florent (2004). Computing the implied volatility in stochastic volatility models. *Communications on Pure and Applied Mathematics 57*(10), 1352–1373.

Carr, P. P. and R. A. Jarrow (1990). The stop-loss start-gain paradox and option valuation: A new decomposition into intrinsic and time value. *The review of financial studies 3*(3), 469–492.

Craig, I. J. D. and A. D. Sneyd (1988). An alternating-direction implicit scheme for parabolic equations with mixed derivatives. *Computers and Mathematics with Applications 16*, 341–350.

Derman, E. and I. Kani (1998). Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility. *International Journal of Theoretical and Applied Finance 1*, 61–110.

Doust, P. (2012). No-arbitrage SABR. *Journal of Computational Finance 15*(3), 3–31.

Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002, September). Managing smile risk. *Wilmott Magazine 1*, 84–108.

Jackel, P. and C. Kahl (2008, March). Hyp hyp hooray. *Wilmott Magazine 34*, 70–81. New version is available at: http://www.christiankahl.com/publications/HypHypHooray.pdf.

Kennedy, J. E., S. Mitra, and D. Pham (2012). On the Approximation of the SABR Model: A Probabilistic Approach. *Applied Mathematical Finance 19*(6), 553–586.

Obloj, J. (2008, May). Fine-tune your smile, correction to Hagan et al. *Wilmott Magazine 35*, 102–104.

Piterbarg, V. (2007, April). Markovian projection for volatility calibration. *Risk Magazine 20*(4), 84–89.

Tavella, D. and C. Randall (2000). *Pricing Financial Instruments, The Finite Difference Method*. John Wiley & Sons, Inc.

Chapter 3

A numerical scheme for perpetual claims

Morten Karlsmark

Abstract

This paper considers a numerical scheme to price perpetual claims. Contrary to claims with a given maturity date which can be priced solving a PDE, perpetual claims can be priced solving an ODE, often a two boundary problem. We apply a simple finite difference scheme to these two-boundary ODEs. Convergence results are given in a number of cases and we test the method on a variety of different financial examples: We calculate the expected exit time of a diffusion from an interval, price perpetual CDOs on households with exploding credit risk, calculate the value of cash when the interest rate can go negative and value perpetual range accruals in the presence of transaction costs.

3.1 Introduction

In continuous time finance a financial claim can normally be priced solving a PDE which reflects the payoff function of the claim and the dynamics of the underlying in the given model, see for example Björk (2004). But there exists claims without a given maturity date these claims are the so called perpetual claims. Since they have no given maturity date the prices are not given as solutions to PDEs, but if our model is time-homogeneous the prices can be expressed as solutions to ODEs. The ODE will be of second order and is often a two-boundary problem. Only very simple cases can be solved analytically, we can though solve the ODE numerically on a finite interval, if we are able to specify appropriate boundary conditions.

Perpetual claims are not a common market product, but they are interesting because the valuation yields information about finite maturity claims and

the underlying model. Moreover some academic questions can be answered by considering perpetual claims, see for example section 3.5.3.

There exist a large literature on the numerical solution of two boundary problems, see for example Keller (1976) and Ascher et al. (1995). We will use a simple finite difference approach and solve our ODE like an implicit step of a PDE solver. This method is described in for example Ascher et al. (1995). We will give convergence results for some important cases and consider the use of the method on different financial problems.

The paper is organized as follows: Section 3.2 presents the method. Section 3.3 considers convergence of the method for linear problems. Some results are well-known but others have to our knowledge not been considered before. Section 3.4 considers the solution of non-linear problems, here we use the theory on Viscosity solutions to show convergence. Then in section 3.5 we apply the method to some examples. These include: finding the expected exit time of a diffusion from an interval, pricing perpetual CDOs on households with exploding credit risk, valuing cash when interest rate can go negative and pricing perpetual range accruals in the presence of transaction costs. Section 3.6 concludes.

3.2 Discretization of linear second order ODEs

Consider a second order ordinary differential equation on the interval (a, b):

$$\frac{1}{2}\sigma(x)^2 f''(x) + \mu(x) f'(x) - k(x) f(x) = g(x) \tag{3.2.1}$$

with a boundary condition in the point a: $k_1 f(a) + k_2 f'(a) = \tilde{a}$ and a boundary condition in the point b: $k_3 f(b) + k_4 f'(b) = \tilde{b}$.

To ease notation later we also define the infinitesimal operator $L(\cdot, \cdot)$ such that

$$L(x, f(x)) = g(x).$$

To solve the ODE numerically we discretize the derivatives using finite differences. It can easily be seen that if a function f is four times differentiable then

$$\begin{aligned}
f'(x) &= \frac{f(x + \Delta x) - f(x - \Delta x)}{2\Delta x} + O(\Delta x^2) \\
f'(x) &= \frac{f(x + \Delta x) - f(x)}{\Delta x} + O(\Delta x) \\
f'(x) &= \frac{f(x) - f(x - \Delta x)}{\Delta x} + O(\Delta x) \\
f''(x) &= \frac{f(x + \Delta x) - 2f(x) + f(x - \Delta x)}{\Delta x^2} + O(\Delta x^2)
\end{aligned}$$

Discretizing in the points $a = x_0, x_1 \ldots, x_{n-1}, x_n = b$ where $x_i - x_{i-1} = \Delta x$, we obtain $n+1$ equations. Equation $0 < i < n$ reads

$$\frac{1}{2}\sigma(x_i)^2 \frac{f(x_i + \Delta x) - 2f(x_i) + f(x_i - \Delta x)}{\Delta x^2} + \mu(x_i)\frac{f(x_i + \Delta x) - f(x_i - \Delta x)}{2\Delta x} - k(x_i)f(x_i) = g(x_i)$$

and for the point a we get

$$k_1 f(a) + k_2 \frac{f(a + \Delta x) - f(a)}{\Delta x} = \tilde{a}$$

A similar expression is obtained for the point b. Note that we use a one sided finite difference to approximate the first order derivative in a and b, so if $k_2 \neq 0$ the discretization will be of order $O(\Delta x)$.

We have then obtained a system of linear equations

$$AF = G \tag{3.2.2}$$

where $G = (\tilde{a}, g(x_1), \ldots, g(x_{n-1}), \tilde{b})$ and F is the approximation to the function f in the points $x_0, \ldots, x_n$.

This matrix A will depend both on the chosen finite difference approach and on the boundary conditions. A will at least be tridiagonal as we need three points to approximate the second order derivative.

$$A = \begin{pmatrix} k_1 - \frac{k_2}{\Delta x} & \frac{k_2}{\Delta x} & \ldots & & \\ a_1 & -b_1 & c_1 & 0 & \ldots \\ \ddots & \ddots & \ddots & \ddots & \ddots \\ \ldots & 0 & a_{n-1} & -b_{n-1} & c_{n-1} \\ & & \ldots & -\frac{k_4}{\Delta x} & k_3 + \frac{k_4}{\Delta x} \end{pmatrix} \tag{3.2.3}$$

where $a_i = \frac{\sigma(x_i)^2}{2\Delta x^2} - \frac{\mu(x_i)}{2\Delta x}$, $b_i = \frac{\sigma(x_i)^2}{\Delta x^2} + k(x_i)$ and $c_i = \frac{\sigma(x_i)^2}{2\Delta x^2} + \frac{\mu(x_i)}{2\Delta x}$.

In financial applications the boundary conditions will be determined by the product we are pricing. If we are pricing a product with a fixed price on the boundaries a rebate for example we will use the fixed price as boundary conditions (Dirichlet boundary conditions), so we set $k_1 = k_3 = 1$ and $k_2 = k_4 = 0$. Contrary to this some claims are defined on an infinite interval and we can only solve the ODE numerically in a finite interval (x_0, x_n). A normal thing to do is then to choose the interval (x_0, x_n) wide enough and assume linearity at the boundaries, $\sigma(x_0) = \sigma(x_n) = 0$ which gives

$$\begin{aligned} k_1 &= -k(x_0),\ k_2 = \mu(x_0),\ G_0 = g(x_0), \\ k_3 &= -k(x_n),\ k_4 = \mu(x_n),\ G_n = g(x_n), \end{aligned} \tag{3.2.4}$$

see Tavella and Randall (2000). Another way to define the boundary conditions is to apply an asymptotic value in both ends, see for example Heider (2010).

It is also possible to use winding ie, approximate the first order derivatives with a single sided finite difference in the direction of the drift, this leads to a tridiagonal matrix A_{wind} where row i will have the tridiagonal elements

$$\begin{array}{l} \frac{\sigma(x_i)^2}{2\Delta x^2} - \frac{\min(\mu(x_k),0)}{\Delta x}, \\ -\frac{\sigma(x_k)^2}{\Delta x^2} + \frac{\min(\mu(x_k),0)}{\Delta x} - \frac{\max(\mu(x_k),0)}{\Delta x} - k(x_k), \\ \frac{\sigma(x_k)^2}{2\Delta x^2} + \frac{\max(\mu(x_k),0)}{\Delta x} \end{array}$$

This finite difference approximation will only be of order $O(\Delta x)$, but for some ODEs it is a very useful discretization, see section 3.3.

Yet another trick that can be used is nonuniform spacing, if we define $\Delta x_i = x_i - x_{i-1}$ we can construct a matrix A_{non} where row i will have the tridiagonal elements:

$$\begin{array}{l} \frac{\sigma(x_i)^2}{\Delta x_i(\Delta x_i + \Delta x_{i+1})} - \frac{\mu(x_i)\Delta x_{i+1}}{\Delta x_i(\Delta x_i + \Delta x_{i+1})}, \\ -\frac{\sigma(x_i)^2}{\Delta x_i \Delta x_{i+1}} + \mu(x_i)\left(\frac{\Delta x_{i+1}}{\Delta x_i(\Delta x_i + \Delta x_{i+1})} - \frac{\Delta x_i}{\Delta x_{i+1}(\Delta x_i + \Delta x_{i+1})}\right) - k(x_i), \\ \frac{\sigma(x_i)^2}{\Delta x_{i+1}(\Delta x_i + \Delta x_{i+1})} + \frac{\mu(x_i)\Delta x_i}{\Delta x_{i+1}(\Delta x_i + \Delta x_{i+1})} \end{array}$$

This finite difference approximation will only be of order $O(\Delta x_{i+1} - \Delta x_i) + O\left(\frac{\Delta x_i^3 + \Delta x_{i+1}^3}{\Delta x_i + \Delta x_{i+1}}\right)$ since the term approximating the second order derivative is of order one. But if $\Delta x_{i+1} - \Delta x_i = O(\Delta x_i^2)$ we still obtain a second order scheme. Often a well chosen nonuniform discretization will outperform a simple uniform discretization.

We have now approximated the ODE with a set of linear equations (3.2.2) that we need to solve. As A in these examples is tridiagonal we can use a simple tridag algorithm to solve the system. If A has more bands we can use LU decomposition. For more on the solution of linear equation systems see Press et al. (2002).

3.3 Convergence of the method

Assume that the true solution to the ODE is four times differentiable. We wish to approximate the solution of (3.2.1) in the points $(x_0, \ldots, x_n)$. By f we denote the true solution in the points $(x_0, \ldots, x_n)$ and by F we denote the

solution to (3.2.2). If we neglect the possible first order discretization of the boundary conditions and apply A to the difference $F - f$ we get

$$A(F - f) = G - G + O(\Delta x^2) = O(\Delta x^2) \tag{3.3.1}$$

as Af approximates the differentials of f to the order Δx^2. This is normally referred to as consistency of the scheme. (3.3.1) implies

$$F - f = A^{-1}O(\Delta x^2)$$

so if every absolute row sum in A^{-1} is $O(1)$ we have $O(\Delta x^2)$ convergence. When every absolute row sum in A^{-1} is $O(1)$ the scheme is said to be stable.

If we use winding to approximate the differentials we get

$$F - f = A_{wind}^{-1}O(\Delta x) \tag{3.3.2}$$

and thereby slower convergence. But the row sums of A_{wind}^{-1} are normally easier to bound than the row sums of A^{-1}.

Let us now turn to different tridiagonal matrices and show that the absolute row sums of the inverse are bounded. The first result is well-known see Ascher et al. (1995) again, but we will give a proof anyways.

Proposition 3.3.1. *Consider the matrix*

$$A = \begin{pmatrix} -b_0 & c_0 & 0 & \dots & & \\ a_1 & -b_1 & c_1 & 0 & \dots & \\ 0 & a_2 & -b_2 & c_2 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \\ & & 0 & a_{n-1} & -b_{n-1} & c_{n-1} \\ & & & 0 & a_n & -b_n \end{pmatrix}$$

If $a_i \geq 0, b_i > 0, c_i \geq 0$ *and* $a_i - b_i + c_i = d_i < K < 0$ *for* $i \in \{1, \dots, n-1\}$, $b_0 > 0, c_0 \geq 0, a_n \geq 0, b_n > 0$, $-b_0 + c_0 = d_0 < K < 0$ *and* $a_n - b_n = d_n < K < 0$, *then* A *is nonsingular and the absolute row sums of* A^{-1} *are bounded by* $\frac{1}{|K|}$.

Proof. $-A$ is clearly a Z-matrix, see definition 2.5.1 in Horn and Johnson (1991). By Gershgorin's circle theorem, see page 31 in Horn and Johnson (1991), we see that all eigenvalues of $-A$ are positive. Therefore $-A$ is an M-matrix, see definition 2.5.2 and 2.1.2 in Horn and Johnson (1991). Using Theorem 2.5.3 in Horn and Johnson (1991) we see that $-A$ is nonsingular and all entries in $-A^{-1}$ are non-negativity so all entries in A^{-1} are non-positive. Also

$$A\ (1, 1, \dots, 1, 1)' = (d_0, \dots, d_n)'$$

which gives

$$(1, 1, \ldots, 1, 1)' = A^{-1} \, (d_0, \ldots, d_n)'$$

As all entries in A^{-1} are non-positive and $(d_0, \ldots d_n)$ is negative and bounded away from zero by K we conclude that the absolute row sums of A^{-1} are bounded by $\frac{1}{|K|}$. □

Example 3.3.2. *Consider an ODE on the form* (3.2.1) *with Dirichlet boundary conditions. We discretize it using a matrix on the form* (3.2.3). *Assume* $k(x) > K > 0$, $\frac{\sigma(x)^2}{2\Delta x^2} - \frac{\mu(x)}{2\Delta x} > 0$ *and* $\frac{\sigma(x)^2}{2\Delta x^2} + \frac{\mu(x)}{2\Delta x} > 0$ *for* $x \in (a, b)$. *It is easily seen that* $d_i = -k(x_i)$, $a_i, b_i, c_i > 0$ *for* $1 \leq i \leq n-1$. *Multiplying equation* 0 *and* n *with* -1 *if necessary we get that* $b_0, b_n > 0$ *and* $c_0 = a_n = 0$. *Using proposition 3.3.1 we see that the row sums of* A^{-1} *are* $O(1)$.

Example 3.3.3. *Consider an ODE on the form* (3.2.1) *and let us apply our linear boundary conditions* (3.2.4). *Assuming the same conditions as in example 3.3.2 and now also that* $\mu(x_0) \geq 0$ *and* $\mu(x_n) \leq 0$ *we get* $b_0, b_n > 0$ *and* $c_0, a_n \geq 0$. *Again we can use proposition 3.3.1.*

Example 3.3.4. *Consider example 3.3.2 but now only assume* $k(x) > K > 0$ *for* $x \in (a, b)$. *If we use the normal finite difference discretization we cannot fulfill the assumptions in proposition 3.3.1. But if we use winding we get*

$$\begin{aligned} a_i &= \frac{\sigma(x_i)^2}{2\Delta x^2} \\ -b_i &= -\frac{\sigma(x_i)^2}{\Delta x^2} - \frac{\mu(x_i)}{\Delta x} - k(x_i) \\ c_i &= \frac{\sigma(x_i)^2}{2\Delta x^2} + \frac{\mu(x_i)}{\Delta x} \end{aligned}$$

when $\mu(x_i) > 0$ *and*

$$\begin{aligned} a_i &= \frac{\sigma(x_i)^2}{2\Delta x^2} - \frac{\mu(x_i)}{\Delta x} \\ -b_i &= -\frac{\sigma(x_i)^2}{\Delta x^2} + \frac{\mu(x_i)}{\Delta x} - k(x_i) \\ c_i &= \frac{\sigma(x_i)^2}{2\Delta x^2} \end{aligned}$$

when $\mu(x_i) < 0$. *We can then apply Proposition 3.3.1 again.*

Note that A_{non} also will fulfill the conditions under similar assumptions.

When $d_i = a_i - b_i + c_i$ cannot be bounded away from zero we need another result. The next proposition gives some conditions under which we still obtain convergence.

Proposition 3.3.5. *Consider a matrix*

$$A = \begin{pmatrix} -b_0 & c_0 & 0 & \dots & & \\ a_1 & -b_1 & c_1 & 0 & \dots & \\ 0 & a_2 & -b_2 & c_2 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \\ & & 0 & a_{n-1} & -b_{n-1} & c_{n-1} \\ & & & 0 & a_n & -b_n \end{pmatrix}$$

Assume $a_i, b_i, c_i > 0$ *for all* i, $a_i - b_i + c_i \leq 0$ *for* $i \in \{1, \dots n-1\}$, $b_0 \geq kc_0$, $b_n \geq ka_n$ *where* $k > 1$ *and independent of* n *and* $\frac{1}{b_0} = O(n^\alpha)$, $\frac{1}{b_n} = O(n^\beta)$. *If*

- $c_i \geq a_i$ *for* $i \in \{1, \dots, n-1\}$ *and* $\max_{i \in \{1,\dots,n-1\}} \frac{1}{c_i} = O(n^{-\gamma})$,

or

- $a_i \geq c_i$ *for* $i \in \{1, \dots, n-1\}$ *and* $\max_{i \in \{1,\dots,n-1\}} \frac{1}{a_i} = O(n^{-\gamma})$.

then A *is nonsingular and the absolute row sums of* A^{-1} *are* $O(n^{2-\gamma} + n^\alpha + n^\beta)$.

Proof. Using J_{30} from page 136 in Berman and Plemmons (1979) on the vector $(1, \dots, 1)^T$ and $-A$ we see that $-A$ is a nonsingular M-matrix, so A is nonsingular.

For $i \in \{0, \dots, n\}$ define

$$\tau_i = \frac{|c_i|}{|b_i| - |a_i|} \quad \omega_i = \frac{|a_i|}{|b_i| - |c_i|}$$

with definition $a_0 = c_n = 0$. Note that both $0 \leq \tau_i \leq 1$ and $0 \leq \omega_i \leq 1$. Then define

$$\tau_{i,0} = \tau_i$$
$$\tau_{i,j} = \begin{cases} \tau_{i,j-1} & i < j \\ \frac{|c_i|}{|b_i| - \tau_{i-1,j-1}|a_i|} & \text{else} \end{cases}$$

$$\omega_{i,0} = \omega_i$$
$$\omega_{i,j} = \begin{cases} \omega_{i,j-1} & n - j < i \\ \frac{|a_i|}{|b_i| - \omega_{i+1,j-1}|c_i|} & \text{else} \end{cases}$$

where $0 \leq \tau_{i,j} \leq 1$ and $0 \leq \omega_{i,j} \leq 1$ by induction.

Let $(m_{ij}) = M = A^{-1}$, theorem 3.4 in Nabben (1999) states that for any $t \in \{0, \dots, n-1\}$ we have

$$\begin{aligned} |m_{i,j}| &\leq |m_{j,j}| \Pi_{k=i}^{j-1} \tau_{k,t} \qquad \text{for } i < j \\ |m_{i,j}| &\leq |m_{j,j}| \Pi_{k=j+1}^{i} \omega_{k,t} \quad \text{for } i > j \end{aligned} \tag{3.3.3}$$

and theorem 3.5 in Nabben (1999) states that

$$\frac{1}{|b_i| + \tau_{i-1,t}|a_i| + \omega_{i+1,t}|c_i|} \leq |m_{i,i}| \leq \frac{1}{|b_i| - \tau_{i-1,t}|a_i| - \omega_{i+1,t}|c_i|}$$

for $t \in \{0, \ldots, n-1\}$.

Actually theorem 3.5 Nabben (1999) also requires that $|b_i| - \tau_{i-1}|a_i| - \omega_{i+1}|c_i| \neq 0$. This seems to be a typo, since the result can be obtained without this assumption.

In order to use the results we need to calculate $\tau_{i,j}$ and $\omega_{i,j}$ for our matrix. Let us consider the case where $c_i \geq a_i$, then we only need to look at $\omega_{i,j}$.

We have

$$\begin{aligned} \omega_{n,0} &\leq \frac{1}{k} \leq \frac{l}{l+1} \quad \text{for some } l \in \mathbb{N} \\ \omega_{i,0} &\leq 1 \quad \text{for} \quad i \in \{1, \ldots, n-1\} \end{aligned}$$

$\omega_{0,0}$ is not needed.

Let us then look at the recursive equation, we calculate $\omega_{i,j}$ when $n-j=i$:

$$\begin{aligned} \omega_{i,j} &= \frac{|a_i|}{|b_i| - \omega_{i+1,j-1}|c_i|} \\ &= \frac{a_i}{b_i - \omega_{i+1,j-1}c_i} \\ &\leq \frac{a_i}{b_i - \omega_{i+1,j-1}(b_i - a_i)} \\ &= \frac{1}{(1 - \omega_{i+1,j-1})\frac{b_i}{a_i} + \omega_{i+1,j-1}} \\ &\leq \frac{1}{2(1 - \omega_{i+1,j-1}) + \omega_{i+1,j-1}} \\ &= \frac{1}{2 - \omega_{i+1,j-1}} \end{aligned}$$

The first inequality holds because $c_i \leq b_i - a_i$, and the last inequality holds because $c_i \geq a_i$. Using this recursive formula and knowing $\omega_{n,0} \leq \frac{l}{l+1}$ we see that

$$\omega_{n-i,n-1} \leq \frac{l+i}{l+i+1} = \frac{l+n-(n-i)}{l+n-(n-i)+1}$$

So we obtain an upper bound for $m_{i,i}$ when $0 < i < n$:

$$\begin{aligned}|m_{i,i}| &\leq \frac{1}{|b_i| - \tau_{i-1,n-1}|a_i| - \omega_{i+1,n-1}|c_i|} \\ &\leq \frac{1}{b_i - a_i - \frac{l+n-i-1}{l+n-i}c_i} \\ &\leq \frac{1}{a_i + c_i - a_i - \frac{l+n-i-1}{l+n-i}c_i} \\ &= \frac{l+n-i}{c_i}\end{aligned}$$

We also have $|m_{0,0}| \leq \frac{1}{b_0 - c_0} \leq \frac{1}{b_0(1-\frac{l}{l+1})} = \frac{l+1}{b_0}$ and $|m_{n,n}| \leq \frac{1}{b_n - a_n} \leq \frac{1}{b_n(1-\frac{l}{l+l})} = \frac{l+1}{b_n}$. Using (3.3.3) we get

$$|m_{i,j}| \leq \frac{l+n-j}{c_j}$$

for $0 < j < n$, $|m_{i,0}| \leq \frac{l+1}{b_0}$ and $|m_{i,n}| \leq \frac{l+1}{b_n}$.

The absolute sum of row i is therefore bounded by

$$\sum_{j=1}^{n-1} \frac{l+n-j}{c_j} + \frac{l+1}{b_0} + \frac{l+1}{b_n} \leq n(l+n) \max_{0<j<n} \frac{1}{c_j} + \frac{l+1}{b_0} + \frac{l+1}{b_n} = O(n^{2-\gamma} + n^{\alpha} + n^{\beta})$$

The proof for the $a_i \geq c_i$ case is parallel, just look at τ instead of ω. □

Proposition 3.3.6. *Consider a matrix*

$$A = \begin{pmatrix} -1 & 0 & 0 & \dots & & \\ a_1 & -b_1 & c_1 & 0 & \dots & \\ 0 & a_2 & -b_2 & c_2 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \\ & & 0 & a_{n-1} & -b_{n-1} & c_{n-1} \\ & & & 0 & 0 & -1 \end{pmatrix}$$

Assume $a_i, b_i, c_i > 0$, $a_i - b_i + c_i \leq 0$, $b_1 \geq kc_1$, $b_{n-1} \geq ka_{n-1}$ where $k > 1$ and independent of n and $\frac{1}{b_1} = O(n^{\alpha})$, $\frac{1}{b_{n-1}} = O(n^{\beta})$. If

- *$c_i \geq a_i$ for $i \in \{1, \dots, n-1\}$ and $\max_{i \in \{1,\dots,n-1\}} \frac{1}{c_i} = O(n^{-\gamma})$,*

or

- *$a_i \geq c_i$ for $i \in \{1, \dots, n-1\}$ and $\max_{i \in \{1,\dots,n-1\}} \frac{1}{a_i} = O(n^{-\gamma})$.*

Then A is nonsingular and the absolute row sums of A^{-1} are $O(n^{2-\gamma} + n^{\alpha} + n^{\beta} + 2)$.

Proof. Note that we cannot use proposition 3.3.5 directly since $c_0 = a_n = 0$.

But as in the proof of proposition 3.3.5 we can use J_{30} from page 136 in Berman and Plemmons (1979) on the vector $(1, \dots, 1)^T$ and $-A$ to see that $-A$ is a nonsingular M-matrix, so A is nonsingular.

Let us then define the matrix

$$B = \begin{pmatrix} -b_1 & c_1 & 0 & \dots & & \\ a_2 & -b_2 & c_2 & 0 & \dots & \\ 0 & a_3 & -b_3 & c_3 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \\ & & \dots & 0 & a_{n-1} & -b_{n-1} \end{pmatrix}$$

so

$$A = \begin{pmatrix} -1 & 0 & \dots & \dots & 0 \\ a_1 & & & & \vdots \\ 0 & & B & & 0 \\ \vdots & & & & c_{n-1} \\ 0 & \dots & & 0 & -1 \end{pmatrix}$$

note that B is nonsingular by proposition 3.3.5

It is easy to verify that

$$A^{-1} = \begin{pmatrix} -1 & 0 & \dots & 0 & 0 \\ v^1 & & B^{-1} & & v^2 \\ 0 & 0 & \dots & 0 & -1 \end{pmatrix}$$

where $v_i^1 = a_1 \left(B^{-1}\right)_{i,0}$ and $v_i^2 = c_{n-1} \left(B^{-1}\right)_{i,n-2}$ for $i \in \{0, \dots, n-2\}$. Using the proof of proposition 3.3.5 we obtain

$$\left|v_i^1\right| = a_1 \left|\left(B^{-1}\right)_{i,0}\right| \leq a_1 \frac{1}{b_1 - c_1} \leq 1$$

and

$$\left|v_i^2\right| = c_{n-1} \left|\left(B^{-1}\right)_{i,n-2}\right| \leq c_{n-1} \frac{1}{b_{n-1} - a_{n-1}} \leq 1$$

for $i \in \{0, \dots, n-2\}$. Which shows that the absolute sum of row $i \in \{1, \dots, n-1\}$ in A^{-1} can be bounded by the absolute sum of row $i-1$ in B^{-1} plus 2. Also the absolute sum of row 0 and n are 1. Using proposition 3.3.5 on B gives us the result. □

Example 3.3.7. *Consider an ODE on the form* (3.2.1) *with Dirichlet boundary conditions. Assume* $k(x) \geq 0$, $\sigma(x) > K > 0$ *and* $0 \leq \mu(x) < K_\mu$ *on* (a, b).

We get $c_i \geq a_i$, $a_i - b_i + c_i \leq 0$, $\frac{1}{c_i} = O(\Delta x^2) = O(n^{-2})$ *for* $i \in \{1, \dots n-1\}$, $\frac{1}{b_1} = O(n^{-2})$ *and* $\frac{1}{b_{n-1}} = O(n^{-2})$. *The Dirichlet boundary conditions gives us* $b_0 = b_n = 1$ *and* $c_0 = a_n = 0$, *and for* n *large enough we have* $a_i, b_i, c_i > 0$, $b_1 > \frac{3}{2}c_1$ *and* $b_{n-1} > \frac{3}{2}a_{n-1}$. *Using proposition 3.3.6 we see that the absolute sum of a row in* A^{-1} *is* $O(1)$. *The case* $K_\mu < \mu(x) \leq 0$ *can be handled in the same way.*

Example 3.3.8. *Consider an ODE on the form* (3.2.1) *given on the interval* $(0, b)$ *with Dirichlet boundary conditions. Assume* $k(x) \geq 0$, $\sigma(x) = x^\beta$, $0 < \beta < 0.5$ *and* $0 \leq \mu(x) < K_\mu$. *We get* $c_i \geq a_i$, $a_i - b_i + c_i \leq 0$, $\frac{1}{c_i} = O(\Delta x^{2-2\beta}) = O(n^{-2(1-\beta)})$, $\frac{1}{b_1} = O(n^{-2(1-\beta)})$ *and* $\frac{1}{b_{n-1}} = O(n^{-2})$. *The Dirichlet boundary conditions gives us* $b_0 = b_n = 1$ *and* $c_0 = a_n = 0$, *and for* n *large enough we have* $a_i, b_i, c_i > 0$, $b_1 > \frac{3}{2}c_1$ *and* $b_{n-1} > \frac{3}{2}a_{n-1}$ *because* $0 < \beta < 0.5$. *Using proposition 3.3.6 we see that the absolute sum of a row in* A^{-1} *is* $O(n^{2-2(1-\beta)}) = O(n^{2\beta})$. *The finite difference scheme will therefore have an error of* $O(n^{2\beta-2})$, *but only when* $\beta < 0.5$.

3.4 Methods for Non-linear ODEs

In some cases the price of a product is described by a non-linear ODE. The ODE will normally be non-linear in the second order derivative or be an obstacle problem. See for example Heider (2010) that deals with the convergence of finite difference schemes for non-linear PDEs. We will solve the non-linear ODEs in the viscosity sense using a result from Barles and Souganidis (1991). For an introduction to Viscosity solutions we refer to Crandall et al. (1992).

Consider a second order non-linear ODE on the interval (a, b)

$$H(x, f(x), f'(x), f''(x)) = \frac{1}{2}\sigma(x, f''(x))^2 f''(x) + \mu(x)f'(x) - k(x)f(x) - g(x) = 0 \tag{3.4.1}$$

with a boundary condition in the point a: $k_1 f(a) + k_2 f'(a) = \tilde{a}$ and a boundary condition in the point b: $k_3 f(b) + k_4 f'(b) = \tilde{b}$. Also assume a viscosity solution exists to the equation.

We discretize the equation using finite differences to obtain a system

$$A(F)F = G \tag{3.4.2}$$

A is tridiagonal and the nonzero elements of row i are

$$\begin{aligned}
&\frac{\sigma\left(x_i, \frac{F_{i+1}-2F_i+F_{i-1}}{\Delta x^2}\right)^2}{2\Delta x^2} - \frac{\mu(x_i)}{2\Delta x}, \\
&-\frac{\sigma\left(x_i, \frac{F_{i+1}-2F_i+F_{i-1}}{\Delta x^2}\right)^2}{\Delta x^2} - k(x_i), \\
&\frac{\sigma\left(x_i, \frac{F_{i+1}-2F_i+F_{i-1}}{\Delta x^2}\right)^2}{2\Delta x^2} + \frac{\mu(x_i)}{2\Delta x}
\end{aligned}$$

note that the matrix A depends on the discrete function F.

In order to show convergence of the scheme to the viscosity solution of (3.4.1) we need to show that the scheme (3.4.2) is monotone, stable, consistent and that a strong comparison result holds for the differential equation (3.4.1). We will from now on assume a strong comparison result holds for the differential equation (3.4.1).

Proposition 3.4.1. *Assume there exists a solution F to the equation system* (3.4.2) *for all $\Delta x < \Delta$ where $\Delta > 0$ is some constant. If*

- *$\delta > 0$ exists such that $\forall \epsilon > 0, x \in (a,b), f$*

$$\sigma\left(x, f+\epsilon\right)^2 \left(f+\epsilon\right) \geq \sigma\left(x, f\right)^2 f + \delta\epsilon \tag{3.4.3}$$

- *$\sigma(x,\xi)^2\xi$ is continuous and $\sigma(\cdot,\cdot)^2 > K > 0$*
- *$\mu(x)$ and $g(x)$ are bounded on (a,b).*
- *We have Dirichlet boundary conditions or linear boundary conditions on the form* (3.2.4) *with $\mu(a) \geq 0$ and $\mu(b) \leq 0$.*
- *$k(x) > K > 0$*

Then the scheme (3.4.2) *converges to the viscosity solution of* (3.4.1).

Proof. Monotonicity:

We will just verify condition (2.2) in Barles and Souganidis (1991). Other authors suggest longer definitions see Pooley et al. (2003) and D'Halluin et al. (2005).

Consider first equation 1 to $n-1$. We perturb F_{i+1} and F_{i-1} in equation i with $\xi_{i+1} > 0$ and $\xi_{i-1} > 0$. Then we subtract the original equation. Thereby

we obtain

$$\begin{aligned}
&\frac{1}{2}\sigma\left(x_i, \frac{F_{i+1}+\xi_{i+1}-2F_i+F_{i-1}+\xi_{i-1}}{\Delta x^2}\right)^2 \frac{F_{i+1}+\xi_{i+1}-2F_i+F_{i-1}+\xi_{i-1}}{\Delta x^2} \\
&\quad + \mu(x_i)\frac{F_{i+1}+\xi_{i+1}-F_{i-1}-\xi_{i-1}}{\Delta x} \\
&\quad - \frac{1}{2}\sigma\left(x_i, \frac{F_{i+1}-2F_i+F_{i-1}}{\Delta x^2}\right)^2 \frac{F_{i+1}-2F_i+F_{i-1}}{\Delta x^2} - \mu(x_i)\frac{F_{i+1}-F_{i-1}}{\Delta x} \\
&\geq \xi_{i+1}\left(\frac{\delta}{2\Delta x^2}+\frac{\mu(x_i)}{\Delta x}\right) + \xi_{i-1}\left(\frac{\delta}{2\Delta x^2}-\frac{\mu(x_i)}{\Delta x}\right)
\end{aligned}$$

This is non-negative for Δx small.

Then we consider the boundary conditions: If we have Dirichlet boundary conditions (multiplied with -1 in order to fit them into the matrix) we get the equations

$$\begin{aligned}
-F_0 &= -\tilde{a} \\
-F_n &= -\tilde{b}
\end{aligned}$$

these will not change when we perturb F_1 and F_{n-1}.

With our linear boundary conditions (3.2.4) we get

$$-k(a)F_0 + \mu(a)\frac{F_1 - F_0}{\Delta x} = g(a) \tag{3.4.4}$$

$$-k(b)F_n + \mu(b)\frac{F_n - F_{n-1}}{\Delta x} = g(b) \tag{3.4.5}$$

A perturbation with $\xi_1 > 0$ in F_1 does not decrease (3.4.4) when $\mu(a) \geq 0$. Similarly a perturbation with $\xi_{n-1} > 0$ in F_{n-1} does not decrease (3.4.5) when $\mu(b) \leq 0$. So the scheme is monotone. Note that we actually have shown the reverse inequality of (2.2) in Barles and Souganidis (1991), we can of course reverse the sign by multiplying the entire system with -1.

Consistency:

Consistency is simply satisfied because $\sigma(x,\xi)^2\xi$ is continuous.

Stability:

Since a solution to $A(F)F = G$ exists when Δx is small, $\mu(x)$ is bounded on (a,b), $k(x) > K > 0$ and $\sigma(\cdot,\cdot)^2 > K > 0$ we can use proposition 3.3.1 to obtain

$$\max_i |F_i| \leq \max_i \frac{|G_i|}{|K|}$$

this implies stability since $g(x)$ is bounded on (a,b).

The result follows by Theorem 2.1 in Barles and Souganidis (1991). □

The system (3.4.2) is normally solved using the Newton method, see for example Ascher et al. (1995). In some models the Newton method can be analyzed in full detail, this is for example true in the uncertain volatility model of Avellaneda et al. (1995) and Lyons (1995), see Pooley et al. (2003).

3.4.1 American options

Next we wish to approximate ODEs for American options. Therefore consider the equation

$$\max\left(H\left(x, f(x), f'(x), f''(x)\right), \phi(x) - f(x)\right) = 0 \tag{3.4.6}$$

with boundary conditions in the points a and b

$$\begin{aligned}\max\left(k_1 f(a) + k_2 f'(a) - \tilde{a}, \phi(a) - f(a)\right) &= 0\\ \max\left(k_3 f(b) + k_4 f'(b) - \tilde{b}, \phi(b) - f(b)\right) &= 0\end{aligned}$$

Here $H(\cdot)$ has the form (3.4.1). Again we assume the existence of a viscosity solution to the differential equation. Theorem 6.7 in Touzi (2010) shows that the value of some optimal stopping problems will be a viscosity solution to (3.4.6).

Approximating $H(\cdot)$ using a matrix $A(F)$ and a vector G we obtain the linear complementarity problem

$$\max\left(A(F)F - G, \Phi - F\right) = 0 \tag{3.4.7}$$

which can be written as

$$\begin{aligned}&A(F)F \leq G\\ &\Phi - F \leq 0\\ &(A(F)F - G)'(\Phi - F) = 0.\end{aligned}$$

Note that we let $A(\cdot)$ depend on F such that we can capture nonlinearities. As before we will assume that a strong comparison principle holds for (3.4.6).

Proposition 3.4.2. *Assume there exists a solution F to the equation system (3.4.7) for all $\Delta x < \Delta$ where $\Delta > 0$ is some constant. If*

- *$\delta > 0$ exists such that $\forall \epsilon > 0, x \in (a, b), f$*

$$\sigma\left(x, f + \epsilon\right)^2 \left(f + \epsilon\right) \geq \sigma\left(x, f\right)^2 f + \delta\epsilon, \tag{3.4.8}$$

- *$\sigma(x, \xi)^2 \xi$ is continuous and $\sigma(\cdot, \cdot)^2 > K > 0$*
- *$\mu(x)$, $g(x)$ and $\phi(x)$ is bounded on (a, b).*

- *We have Dirichlet boundary conditions or our linear boundary conditions* (3.2.4) *with* $\mu(a) \geq 0$ *and* $\mu(b) \leq 0$.
- $k(x) > K > 0$

Then the scheme (3.4.7) *converge to the viscosity solution of* (3.4.6).

Proof. Monotonicity:

Consider equation $i \in \{1, \ldots, n-1\}$ in (3.4.7) and perturb F_{i-1} with $\xi_{i-1} > 0$ and F_{i+1} with $\xi_{i+1} > 0$. From the proof of proposition 3.4.1 we see that the first part of (3.4.7) does not decrease for Δx small and the second part stays the same. Using the same arguments we also see that equation 0 and n are monotone. Therefore we conclude that the scheme is monotone.

Consistency:

Consistency follows since $A(F)F - G$ is consistent for the expression $H\left(x, f(x), f'(x), f''(x)\right)$.

Stability:

The solution to (3.4.7) must satisfy

$$\tilde{A}(F)F = \tilde{G} \tag{3.4.9}$$

where the matrix $\tilde{A}(F)$ is defined by

$$\begin{aligned} \tilde{A}(F)_{i,j} &= A(F)_{i,j} \quad \forall j \qquad \text{when} \quad (A(F)F)_i = G_i \\ \tilde{A}(F)_{i,j} &= \begin{cases} 0 & j \neq i \\ -1 & j = i \end{cases} \qquad \text{when} \quad F_i = \Phi_i \end{aligned}$$

and the vector $\tilde{G}$ by

$$\begin{aligned} \tilde{G}_i &= G_i \qquad \text{when} \quad (A(F)F)_i = G_i \\ \tilde{G}_i &= -\Phi_i \qquad \text{when} \quad F_i = \Phi_i \end{aligned}$$

Using proposition 3.3.1 on the system (3.4.9) we see that

$$\max_i |F_i| \leq \max_i \frac{|\tilde{G}_i|}{\min(|K|, 1)}$$

So F is bounded since $g(\cdot)$ and $\phi(\cdot)$ are bounded on (a, b). Again the result follows by Theorem 2.1 in Barles and Souganidis (1991). □

If $A(F)F$ in (3.4.7) is linear ($A(F) = A$) and $-A$ is a P-matrix, we can solve the problem using the PSOR algorithm, see Tavella and Randall (2000). If $-A$ also is a tridiagonal M-matrix we can use the Cryer algorithm, see Cryer (1983).

3.5 Applications

There exist a large number of applications for the methods considered here. We will present four numerical examples. As a simple first example we calculate the expected exit time of a diffusion from a certain interval. The second example will be the valuation of perpetual CDO tranches on households with explosive credit risk. The third will be the value of cash in an environment where interest rates can be negative. The last example will consider the pricing of a perpetual range accrual when transaction costs are present.

3.5.1 Exit times

We wish to find the expected time when a stochastic process leaves a given interval. So we like to find the expected exit time. Assume we have a process

$$dX_t = \mu(X_t)\mathrm{d}t + \sigma(X_t)\mathrm{d}W_t \tag{3.5.1}$$

and an interval (a, b). Let $\tau = \inf\{t \geq 0 : X_t \notin (a, b)\}$ (the exit time) and define

$$f(x) = E(\tau - t | X_t = x)$$

We must have $f(a) = f(b) = 0$ and

$$f(X_0) = E(\tau) = E\left(\int_0^\tau 1 \ \mathrm{d}s\right) \tag{3.5.2}$$

Using Ito's lemma on f we obtain

$$f(X_\tau) = f(X_0) + \int_0^\tau \mu(X_s)f'(X_s) + \frac{1}{2}\sigma(X_s)^2 f''(X_s)\mathrm{d}s + \int_0^\tau \sigma(X_s)f'(X_s)\mathrm{d}W_s$$

where $f(X_\tau) = 0$. If we assume the stochastic integral is a martingale and apply the expectation operator we get

$$f(X_0) = -E\left(\int_0^\tau \mu(X_s)f'(X_s) + \frac{1}{2}\sigma(X_s)^2 f''(X_s)\mathrm{d}s\right) \tag{3.5.3}$$

Comparing (3.5.2) and (3.5.3) we see that if

$$\mu(x)f'(x) + \frac{1}{2}\sigma(x)^2 f''(x) = -1$$

then $f(x)$ is the expected exit time, see also Karatzas and Shreve (1998). This is a two-boundary ODE and example 3.3.7 or 3.3.8 ensures that our finite difference method can solve the ODE if $\mu(x)$ does not change sign on (a, b).

If we also assume that the process is killed with a function $k(x)$, we can find the expected exit or killing time. We simply use the same arguments as

above but we also take a jump into account that brings X_t to the killed state c. So the dynamics of X_t is given by

$$\mathrm{d}X_t = \mu(X_t)\mathrm{d}t + \sigma(X_t)\mathrm{d}W_t + (c - X_t)\mathrm{d}N_t$$

where N_t is a Poisson process with intensity $k(X_t)$. Using Ito's lemma on a function f yields

$$\begin{aligned} f(X_\tau) &= f(X_0) + \int_0^\tau \mu(X_s)f'(X_s) + \frac{1}{2}\sigma(X_s)^2 f''(X_s)\mathrm{d}s + \int_0^\tau \sigma(X_s)f'(X_s)\mathrm{d}W_s \\ &\quad + \int_0^\tau f(c) - f(X_s)\mathrm{d}N_s \\ &= f(X_0) + \int_0^\tau \mu(X_s)f'(X_s) + \frac{1}{2}\sigma(X_s)^2 f''(X_s)\mathrm{d}s + \int_0^\tau \sigma(X_s)f'(X_s)\mathrm{d}W_s \\ &\quad + \int_0^\tau f(c) - f(X_s)\mathrm{d}N_s - \int_0^\tau \left(f(c) - f(X_s)\right) k(X_s)\mathrm{d}s \\ &\quad + \int_0^\tau \left(f(c) - f(X_s)\right) k(X_s)\mathrm{d}s \end{aligned}$$

Taking expectation we obtain

$$\begin{aligned} f(X_0) = -E\Bigg(&\int_0^\tau \mu(X_s)f'(X_s) + \frac{1}{2}\sigma(X_s)^2 f''(X_s)\mathrm{d}s \\ &+ \int_0^\tau \left(f(c) - f(X_s)\right) k(X_s)\mathrm{d}s \Bigg) \end{aligned}$$

and as $f(c) = 0$ we get

$$f(X_0) = -E\left(\int_0^\tau \mu(X_s)f'(X_s) + \frac{1}{2}\sigma(X_s)^2 f''(X_s)\mathrm{d}s - \int_0^\tau f(X_s)k(X_s)\mathrm{d}s\right) \tag{3.5.4}$$

Comparing (3.5.2) and (3.5.4) we see that if

$$\mu(x)f'(x) + \frac{1}{2}\sigma(x)^2 f''(x) - k(x)f(x) = -1$$

then $f(x)$ is the expected exit or killing time.

Example 3.3.2 or 3.3.4 ensures that our numerical method can solve the ODE if $k(x) > K > 0$.

3.5.1.1 Numerical example

We consider the SDE

$$\mathrm{d}X_t = rX_t + \sigma X_t^\gamma \mathrm{d}W_t \tag{3.5.5}$$

assume $\sigma = 0.2$, $r = 0.05$ and let γ vary. We set $a = 1$ and $b = 10$.

In order to compute the expected exit time we will use a nonuniform grid. Heuristically, we define $g(x) = x^{1-\gamma}$ and sees that

$$\mathrm{d}g(X_t) = \ldots \mathrm{d}t + C\mathrm{d}W_t$$

where C is a constant. As the $\mathrm{d}W_t$ part is a Gaussian process we choose to set up an equidistant grid between $g(a)$ and $g(b)$ and transform all the points back with $g^{-1}(\cdot)$, this will then be our nonuniform grid.

Figure 3.1 shows the expected exit time for different starting values X_0. In figure 3.2 we consider the same example but now assume the process is killed with the function

$$k(X_t) = \frac{1}{X_t} \tag{3.5.6}$$

We have used 100 grid points for all the calculations.

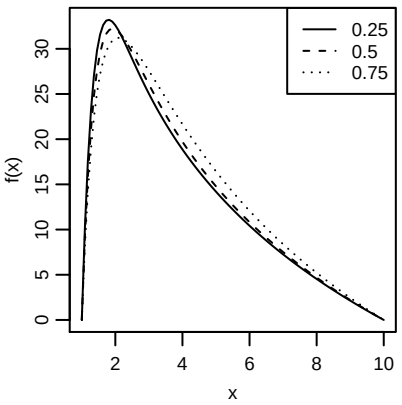

Figure 3.1: Expected exit time from the interval $(1, 10)$ for the process (3.5.5), with parameter values: $\sigma = 0.2$, $r = 0.05$ and γ is varying see the legend.

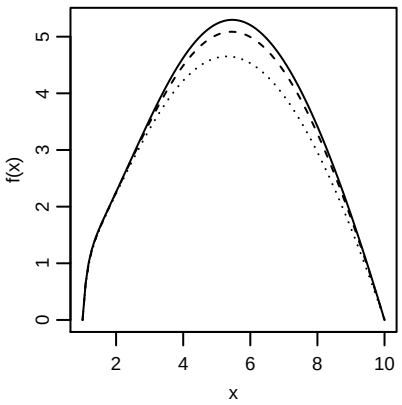

Figure 3.2: Expected exit/killing time from the interval $(1, 10)$ for the process (3.5.5), with parameter values: $\sigma = 0.2$, $r = 0.05$, γ is varying see the legend in figure 3.1. The process is killed with the function (3.5.6).

3.5.2 Perpetual CDOs on households with explosive credit risk

We consider K households and K Poisson processes determining the default times of the households. When the i'th Poisson process jumps the i'th household defaults. The intensity of each Poisson process is given by the process $(\lambda_t)_{t\geq 0}$ and we assume that the Poisson processes are independent conditional on the intensity process. We define

$$\mathrm{d}X_t = \kappa(\theta - X_t)\mathrm{d}t + \sigma X_t^{\gamma}\mathrm{d}W_t$$

and inspired by Andreasen (2001) we let $\lambda_t = \frac{1}{X_t}$. This means that the intensity process can explode in finite time because X_t can hit zero when $\gamma < \frac{1}{2}$ or when $\gamma = \frac{1}{2}$ and $\sigma^2 > 2\kappa\theta$.

Now, let us look at a perpetual CDO contract on the debt of the households. All households pay $a(X_t)\mathrm{d}t$ dollars in a given time interval $\mathrm{d}t$, so the total discounted payment from one household is

$$\int_0^\tau \exp(-rt)a(X_t)\mathrm{d}t$$

where τ is the default time of the household.

We would like to price different tranches of the CDO. A tranche gives the right to a specific part of the payment stream from the households. The payment is determined by the number of defaults in the group of households but different tranches are affected differently by defaults. Let us consider a simple example. We let $K = 10$ and tranche the CDO in three parts: the equity, junior and senior tranche. The equity tranche gives the right to the payments from the first household to default. The junior tranche gives the right to the payments from the next two and the senior tranche gives the right to the payments from the last seven. This means that when no households have defaulted the holder of the equity tranche receives $a(X_t)\mathrm{d}t$, the holder of the junior tranche receives $2a(X_t)\mathrm{d}t$ and the holder of the senior tranche receives $7a(X_t)\mathrm{d}t$. If one household defaults the equity holder will only receive $0\mathrm{d}t$ going forwards but the holders of the junior and senior tranches will be unaffected. If two defaults have happened the junior tranche holder will only receive $a(X_t)\mathrm{d}t$ but the senior tranche holder will still be unaffected etc.

In order to price the different tranches we look at the households in the sequence given by the time of default: The first household to default, the second household to default etc. This we do even though we don't know which household defaults first, but we can specify the default intensity for the first household to default, and thereby value the CDO. This is a consequence of the Poisson processes being independent conditional on the intensity process. The process indicating default of the first household to default is the default processes for the individual households summed up. A sum of independent Poisson processes is a Poisson process with intensities summed up, therefore the process indicating default of the first household to default will be a Poisson process with intensity $K\lambda(t)$. After the first default the process indicating the second default will have intensity $(K-1)\lambda(t)$ etc.

Consider the payment stream from the first household to default. The value at time t is

$$V_t = E_t\left(\int_t^\tau a(X_s)\exp\left(-r(s-t)\right)\mathrm{d}s\right)1_{\{t\leq\tau\}} \tag{3.5.7}$$

where τ now denotes the default time of the first household to default. As V_t

is the value of a perpetual claim the dynamics are

$$\mathrm{d}V_t = \left(V'(X_t)\mathrm{d}X_t + \frac{1}{2}V''(X_t)(\mathrm{d}X_t)^2 + (0 - V(X_t))\mathrm{d}N_t \right) 1_{\{t \leq \tau\}}$$

where N_t is the Poisson process indicating the first default.

Look at the process

$$M_t = \exp(-rt)\, V_t + \int_0^t a(X_s) \exp(-rs)\, 1_{\{s \leq \tau\}} \mathrm{d}s \tag{3.5.8}$$

and assume it is a martingale, which is true if $a(x)$ is bounded. Applying Ito's lemma we get

$$\mathrm{d}M_t = \exp(-rt)\left(-rV_t\mathrm{d}t + a(X_t)\mathrm{d}t + \mathrm{d}V_t\right) 1_{\{t \leq \tau\}}$$

As M_t is a martingale the drift term plus the mean of the Poisson process must be zero, we therefore obtain the differential equation

$$\begin{aligned} 0 &= -rV(x) + a(x) + \kappa(\theta - x)V'(x) + \frac{1}{2}\sigma^2 x^{2\gamma} V''(x) + \frac{K}{x}(0 - V(x)) \Leftrightarrow \\ -a(x) &= -\left(r + \frac{K}{x}\right) V(x) + \kappa(\theta - x)V'(x) + \frac{1}{2}\sigma^2 x^{2\gamma} V''(x) \end{aligned} \tag{3.5.9}$$

Let $V_{i,j}(x)$ denote the value of the payment stream from the i'th household going default given j households have defaulted. $V_{i,i-1}(x)$ is then given by the solution to (3.5.9) if we replace $\frac{K}{x}$ with $\frac{K-(i-1)}{x}$, as we only have $K - (i-1)$ households left. $V_{i,i-2}(x)$ is the solution to

$$\begin{aligned} 0 &= -rV_{i,i-2}(x) + a(x) + \kappa(\theta - x)V'_{i,i-2}(x) + \frac{1}{2}\sigma^2 x^{2\gamma} V''_{i,i-2}(x) \\ &\quad + \frac{K - (i-2)}{x}(V_{i,i-1}(x) - V_{i,i-2}(x)) \Leftrightarrow \\ -a(x) &- \frac{K - (i-2)}{x} V_{i,i-1}(x) = \\ &\quad -\left(r + \frac{K - (i-2)}{x}\right) V_{i,i-2}(x) + \kappa(\theta - x)V'_{i,i-2}(x) + \frac{1}{2}\sigma^2 x^{2\gamma} V''_{i,i-2}(x) \end{aligned}$$

Note the jump condition, when the next default happens we jump to the value where $i-1$ defaults have happened. Continuing this recursion we can easily value the payment stream from the i'th household going default when 0 defaults have happened.

All the differential equations can be solved using our finite difference technique. We use the boundary condition $V(0) = 0$, when the intensity is infinite every household defaults immediately. In the other end we set the second order derivative equal to zero. Doing this we obtain a system on the form

$$AV_{i,j} = G \tag{3.5.10}$$

where $G = \left(0, -a(x_1) - \frac{K-j}{x} V_{i,j+1}(x_1), \ldots, -a(x_n) - \frac{K-j}{x} V_{i,j+1}(x_n)\right)'$ and

$$A = \begin{pmatrix} -1 & 0 & 0 & \ldots & & \\ a_1 & -b_1 & c_1 & 0 & \ldots & \\ \ddots & \ddots & \ddots & \ddots & \ddots & \\ & & \ldots & 0 & -\frac{\kappa(\theta - x_n)}{\Delta x} & \frac{\kappa(\theta - x_n)}{\Delta x} - \frac{K-j}{x_n} - r \end{pmatrix} \tag{3.5.11}$$

where $a_i = \frac{\sigma^2 x_i^{2\gamma}}{2\Delta x^2} - \frac{\kappa(\theta - x_i)}{2\Delta x}$, $b_i = \frac{\sigma^2 x_i^{2\gamma}}{\Delta x^2} + \frac{K-j}{x_i} + r$ and $c_i = \frac{\sigma^2 x_i^{2\gamma}}{2\Delta x^2} + \frac{\kappa(\theta - x_i)}{2\Delta x}$.

As an example let us consider the valuation of our senior tranche. We specify the payoff when different number of households have defaulted. Since the value is zero when all households have defaulted we do not include this in our specification. The payoff specification is then a vector $pay = (1, 2, 3, 4, 5, 6, 7, 7, 7, 7)'$

To value the tranche we start by solving

$$-a(x)pay_1 = -\left(r + \frac{1}{x}\right) V_1(x) + \kappa(\theta - x) V_1'(x) + \frac{1}{2}\sigma x^{2\gamma} V_1''(x)$$

note that the value jumps to zero when default happens. Then for $i = 2$ to $i = 10$ we recursively solve

$$-a(x)pay_i - \frac{i}{x} V_{i-1}(x) = -\left(r + \frac{i}{x}\right) V_i(x) + \kappa(\theta - x) V_i'(x) + \frac{1}{2}\sigma^2 x^{2\gamma} V_i''(x)$$

V_{10} is the value when zero households have defaulted.

To value the junior tranche we specify $pay = (0, \ldots, 0, 1, 2, 2)$ and use the same method. We can of course start with V_8 as $V_7 = 0$. It should now be clear how easy it is to value a perpetual CDO tranche using this method.

We could also include different effects influencing the default of households: Guiso et al. (2013) investigates attitudes towards strategic default on mortgages. You default strategically if you choose to default when the value of your mortgage exceeds the value of your house and you are able to pay the monthly mortgage payment. Guiso et al. (2013) finds that "people who know somebody who defaulted strategically are more likely to declare their intention to do so". This effect could be captured in our model if we let the intensity of the i'th household to default be $f(i)\lambda_t$ instead of $(K - (i-1))\lambda_t$, here $f(1) = K$, $f'(i) \geq -1$ and $f''(i) \geq 0$. When a household defaults it increases the probability that other households will do the same thing.

3.5.2.1 Numerical example

We consider the same example as above so we look at ten households. The equity tranche gives the right to the payments from the first household to default. The junior tranche gives the right to payments from the next two and the senior tranche gives the right to the last seven.

If explosion happens all households default. Therefore the probability of explosion especially influence the price of the senior tranche. By controlling the coefficient γ we control the explosion probability and therefore the price of the senior tranche. Consider the parameters: $r = 0.05$, $\kappa = 0.04$, $\theta = 40$, $\sigma = 5 \cdot 40^{0.5-\gamma}$ and let γ vary. The σ-parameterization is done in order to make everything comparable. For simplicity we let a be a constant but choose it in a way such that the results are on the same scale. For the equity tranche $a = 1$, for the junior tranche $a = \frac{1}{2}$ and for the senior tranche $a = \frac{1}{7}$. We choose the left endpoint of the grid to be 0 and the right endpoint to be $\theta + 5\sqrt{\frac{\theta^{2\gamma}\sigma^2}{2\kappa}}$. Here $\sqrt{\frac{\theta^{2\gamma}\sigma^2}{2\kappa}}$ is a proxy for the stationary standard deviation which is correct for the Vasicek and CIR model. The grid is chosen to be uniform with 500 points, a nonuniform grid does not increase precision. In figure 3.3 to 3.8 we have graphed prices of the different tranches. We have also included the prices from a more normal model where

$$\mathrm{d}\lambda_t = \kappa(\theta - \lambda_t)\mathrm{d}t + \sigma\lambda_t^{\gamma}\mathrm{d}W_t$$

For this model we have chosen the parameters $\kappa = 0.05$, $\theta = 0.07$ and $\sigma = 0.06 \cdot 0.07^{0.5-\gamma}$. This is done in order to approximate the equity tranche price from the explosive model above. We choose a non-uniform grid constructed as in section 3.5.1.1 with 200 points. The left endpoint is $\max\left(\theta - 5\sqrt{\frac{\theta^{2\gamma}\sigma^2}{2\kappa}}, 0\right)$ and the right endpoint is $\theta + 5\sqrt{\frac{\theta^{2\gamma}\sigma^2}{2\kappa}}$. On the boundaries we set the second order derivative equal to zero. See also Ekström et al. (2009) that deals with boundary conditions for term structure equations.

As seen different γ generates large differences in the prices for the senior tranche in our explosive model. Notice the clear difference between the explosive λ model and the non-explosive λ model. The non-explosive parameterization does not provide large price difference for the senior tranche even though the equity tranche is priced differently for different γ. So the explosive model has a nice handle on the senior tranche via the γ parameter.

3.5.3 The value of cash when interest rates can go negative

We follow Andreasen and Bang (2007). Today central banks are obliged to issue coins and notes when a bank account holder demands it. This should imply positive interest rates, since we can choose to hold cash instead of keeping it in the bank. (Holding cash does also imply a cost, insurance for example, which explains the negative interest rates seen in Switzerland recently, we will though neglect these costs). If the central banks were not obliged to issue cash, the interest rate could go negative. A situation that might occur in the future if society thinks a cash system is too expensive to maintain. In this situation cash can be worth more than money on a bank account.

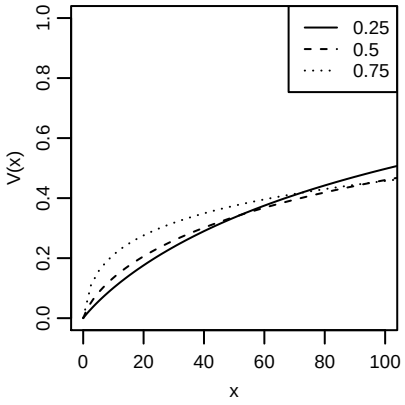

Figure 3.3: The value of the senior tranche of our CDO when the intensity process is explosive. Parameter values: $a = \frac{1}{7}$, $r = 0.05$, $\kappa = 0.05$, $\theta = 40$, $\sigma = 5 \cdot 40^{0.5-\gamma}$ and γ is varying see the legend.

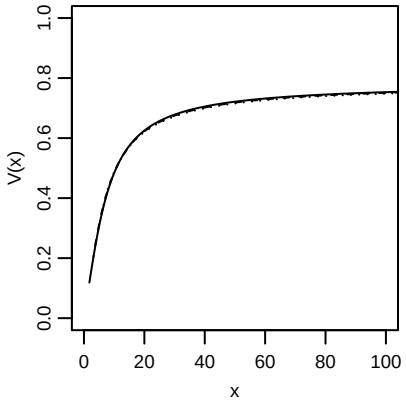

Figure 3.4: The value of the senior tranche of our CDO when the process is non-explosive. Parameter values: $a = \frac{1}{7}$, $r = 0.05$, $\kappa = 0.05$, $\theta = 0.07$, $\sigma = 0.06 \cdot 0.07^{0.5-\gamma}$ and γ is varying see the legend in figure 3.3. Note that $x = \frac{1}{\lambda}$.

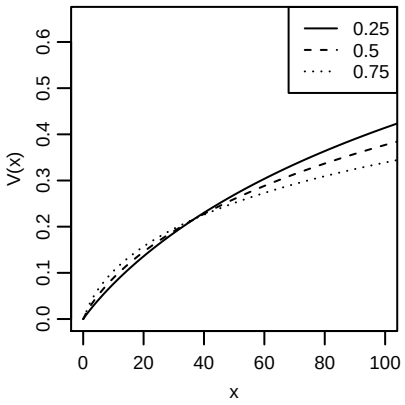

Figure 3.5: The value of the junior tranche of our CDO when the intensity process is explosive. Parameter values: $a = \frac{1}{2}$, $r = 0.05$, $\kappa = 0.05$, $\theta = 40$, $\sigma = 5 \cdot 40^{0.5-\gamma}$ and γ is varying see the legend.

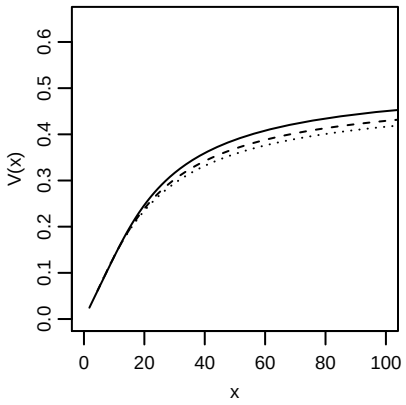

Figure 3.6: The value of the junior tranche of our CDO when the process is non-explosive. Parameter values: $a = \frac{1}{2}$, $r = 0.05$, $\kappa = 0.05$, $\theta = 0.07$, $\sigma = 0.06 \cdot 0.07^{0.5-\gamma}$ and γ is varying see the legend in figure 3.5. Note that $x = \frac{1}{\lambda}$.

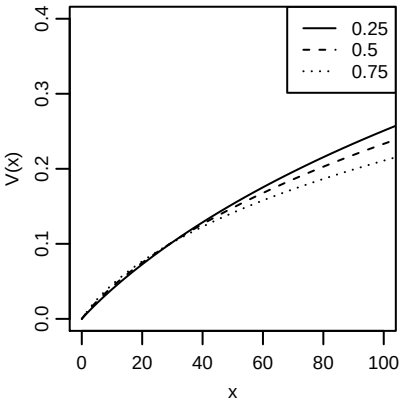

Figure 3.7: The value of the equity tranche of our CDO when the intensity process is explosive. Parameter values: $a = 1$, $r = 0.05$, $\kappa = 0.05$, $\theta = 40$, $\sigma = 5 \cdot 40^{0.5-\gamma}$ and γ is varying see the legend.

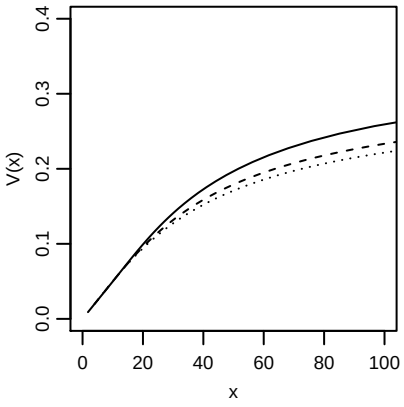

Figure 3.8: The value of the equity tranche of our CDO when the process is non-explosive. Parameter values: $a = 1$, $r = 0.05$, $\kappa = 0.05$, $\theta = 0.07$, $\sigma = 0.06 \cdot 0.07^{0.5-\gamma}$ and γ is varying see the legend in figure 3.7. Note that $x = \frac{1}{\lambda}$.

Let us consider a person with some cash in hand. He can always place the money in the bank and obtain the short rate r_t. When he has done that he cannot get the cash back since it is in fixed supply. So he has to choose the optimal time to place the money in the bank. The value of cash (given in terms of bank account money) is therefore given as a perpetual stopping problem

$$\sup_{\tau} \mathbb{E}^Q \left(\exp\left(- \int_0^\tau r_s \mathrm{d}s \right) \right)$$

Let us then assume

$$\mathrm{d}r_t = \kappa(\theta - r_t)\mathrm{d}t + \sigma \mathrm{d}W_t^Q$$

Using Theorem 6.7 from Touzi (2010) we see that the solution to the stopping problem is a viscosity solution f to the variational inequalities

$$\begin{aligned} 0 &\geq \kappa(\theta - r)f'(r) + \frac{1}{2}\sigma^2 f''(r) - rf(r) \\ 0 &\geq 1 - f(r) \\ 0 &= \left(\kappa(\theta - r)f'(r) + \frac{1}{2}\sigma^2 f''(r) - rf(r) \right)(1 - f(r)) \end{aligned} \tag{3.5.12}$$

In fact the theorem needs r_t to be bounded from below which clearly is violated, we will though use it anyways. Note also that we get a viscosity solution, we

will from now on assume it is the unique one. Other solution methods where one uses the variational inequalities and searches for the maximum solution will yield the same results in the numerical examples below.

We approximate the system using finite differences, and obtain a linear complementarity problem:

$$\begin{aligned} 0 &\geq AF \\ 0 &\geq 1 - F \\ 0 &= (1 - F)' AF \end{aligned} \tag{3.5.13}$$

Here we use the linear boundary conditions for A. Note that the matrix A is not diagonally dominant since r can be negative. Therefore proposition 3.3.1 cannot guarantee stability. Let us instead assume stability and also that a strong comparison principle holds for (3.5.12), then proposition 3.4.2 ensures convergence of the discrete method to the Viscosity solution of (3.5.12).

We solve (3.5.13) using the Cryer algorithm, see Cryer (1983). The algorithm is developed for M-matrices. Computer experiments, where we simply calculate the eigenvalues of $-A$, indicate that for n large enough $-A$ will be an M-matrix in the numerical examples below.

3.5.3.1 Numerical example

We consider the same example as Andreasen and Bang (2007). They calibrate the Vasicek model to the interest rate market in 2007 and obtain the parameters: $\kappa = 3.3\%$, $\theta = 7.7\%$, $\sigma = 0.67\%$. In the following we will denote the optimal stopping short rate by r^*.

Our result is given in figure 3.9. We get $r^* = 0.83\%$, which fits well with the results in Andreasen and Bang (2007). We have used a uniform grid with 5000 points. The endpoints are $\theta \pm 5\sqrt{\frac{\sigma^2}{2\kappa}}$. The large number of points is needed in order to hit r^* precisely.

This result means that you should keep your money in your pocket if the short rate is below 0.83% and if cash is in fixed supply.

We can then investigate what happens when we change σ. Using a uniform grid with 500 grid points and the same endpoints as above we obtain figure 3.10. As seen a larger σ means exercise at a higher short rate, simply because the probability of getting into the negative rate region is higher.

Let us then calibrate the Vasicek model to EUR market data from September 2012. We obtain the parameters $\kappa = 7.1\%$, $\theta = 4.7\%$ and $\sigma = 0.72\%$. The calibration can be done in a number of ways, we have calibrated κ and θ to the yield curve and σ to ATM caplets. Using a uniform grid with 5000 points and the endpoints $\theta \pm 5\sqrt{\frac{\sigma^2}{2\kappa}}$ we get $r^* = 0.68\%$ and the price of cash can be seen in figure 3.11. Again the large number of points is needed in order to hit r^* precisely. Even though θ has decreased and σ has increased the optimal r^* have shifted downwards because κ have gone up.

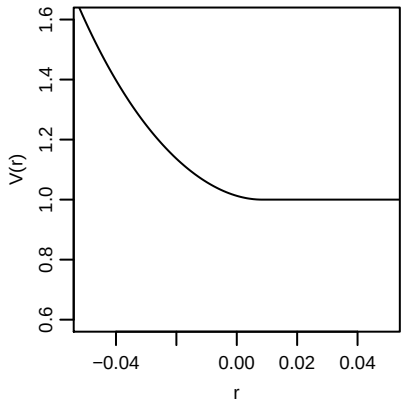

Figure 3.9: Value of cash given in terms of bank account money. Parameters are $\kappa = 3.3\%$, $\theta = 7.7\%$ and $\sigma = 0.67\%$.

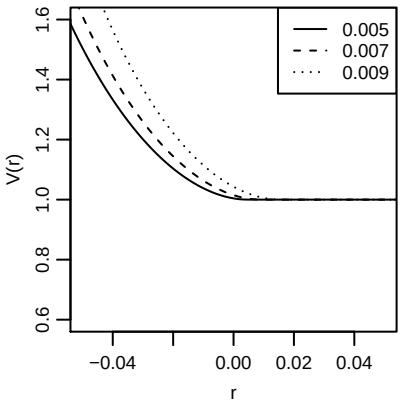

Figure 3.10: Value of cash given in terms of bank account money. Parameters are $\kappa = 3.3\%$, $\theta = 7.7\%$, and σ is varying see the legend.

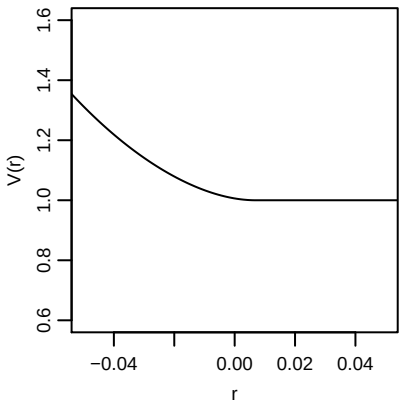

Figure 3.11: Value of cash given in terms of bank account money using 2012 parameters. The parameters are $\kappa = 7.1\%$, $\theta = 4.7\%$ and $\sigma = 0.72\%$

Instead of solving (3.5.13) one can solve a system taking into account continuity of the first derivative. This gives the same result as solving (3.5.13) as it should if smooth fit is a condition for optimality.

3.5.4 Perpetual claims with transaction costs

We will now consider the problem of valuing options taking transaction costs when delta hedging into account.

We consider a stock with price dynamics given by

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma(S_t)\mathrm{d}W_t$$

where $\mu > 0$.

Let us assume we have bought or sold a derivative on the stock, the derivative pays an infinite stream of coupons $g(S_t)\mathrm{d}t$. We delta-hedge this product discretely with a frequency of Δt. Each time we adjust the delta hedge a transaction cost has to be payed. Therefore we have to charge an extra premium on top of the normal price. Leland (1985) develops a model where delta-hedging is done discretely, we will adopt his framework, see also Wilmott et al. (1993) which we follow in our heuristic derivations.

Assume that the transaction costs is a function of the number of stocks we need to trade and the price of the stock at the point of re-hedging: $f(S_t, v_t)$, where v_t is the traded amount.

Consider a portfolio Π consisting of the option position (here we assume the option has been sold) and a delta hedge D:

$$\Pi = -V + DS$$

After a period of time Δt the portfolio has evolved like

$$\Delta\Pi = -\Delta V - g(S)\Delta t + D\Delta S - f(S + \Delta S, \Delta D)$$

where the last term is the transaction cost. Doing a Taylor expansion using Ito's lemma we find that

$$\begin{aligned}\Delta\Pi = \sigma(S)\left(-\frac{\partial V}{\partial s}(S) + D\right)\sqrt{\Delta t}Z \\ + \left(-\frac{1}{2}\sigma(S)^2\frac{\partial^2 V}{\partial s^2}(S) - \mu S\frac{\partial V}{\partial s}(S) + \mu SD - g(S)\right)\Delta t - f(S + \Delta S, \Delta D)\end{aligned}$$

where $Z \sim N(0,1)$. As $D = \frac{\partial V}{\partial S}$ we get

$$\Delta\Pi = -\frac{1}{2}\sigma(S)^2\frac{\partial^2 V}{\partial s^2}(S)\Delta t - g(S)\Delta t - f(S + \Delta S, \Delta D)$$

Again by a Taylor expansion

$$\Delta D = \frac{\partial^2 V}{\partial s^2}(S_t)\sigma(S_t)\sqrt{\Delta t}Z + \frac{\partial^2 V}{\partial s^2}(S_t)rS_t\Delta t + \frac{1}{2}\frac{\partial^3 V}{\partial s^3}(S_t)\sigma(S_t)^2\Delta t + O(\Delta t^{3/2})$$

where $Z \sim N(0,1)$ and independent of $\mathcal{F}_t$.

Then assume $f(s,v) = ks|v|$, the transaction costs is proportional to the stock price and the number of stocks you need to trade. We then have

$$\begin{aligned}
&f(S_{t+\Delta t}, \Delta D) \\
&= kS_{t+\Delta t}\,|\Delta D| \\
&= k\left(S_t + \sigma(S_t)\sqrt{\Delta t}Z + O(\Delta t)\right) \\
&\quad \cdot \left|\frac{\partial^2 V}{\partial s^2}(S_t)\sigma(S_t)\sqrt{\Delta t}Z + \frac{\partial^2 V}{\partial s^2}(S_t)rS_t\Delta t + \frac{1}{2}\frac{\partial^3 V}{\partial s^3}(S_t)\sigma(S_t)^2\Delta t + O(\Delta t^{3/2})\right| \\
&= kS_t\sigma(S_t)\sqrt{\Delta t}\left|\frac{\partial^2 V}{\partial s^2}(S_t)Z\right| + O(\Delta t^{3/2})
\end{aligned}$$

Here we have assumed that

$$\frac{\partial^2 V}{\partial s^2}(S_t)\sigma(S_t)^2 = O(\Delta t^{1/2}), \frac{\partial^2 V}{\partial s^2}(S_t)S_t^2 = O(\Delta t^{1/2}) \quad \text{and}$$
$$\frac{\partial^3 V}{\partial s^3}(S_t) = O(\Delta t^{1/2})$$

which Leland (1985) show is true in his Black-Scholes model. This implies

$$\Delta\Pi = -\frac{1}{2}\sigma(S_t)^2\frac{\partial^2 V}{\partial s^2}(S_t)\Delta t - g(S_t)\Delta t - kS_t\sigma(S_t)\sqrt{\Delta t}\left|\frac{\partial^2 V}{\partial S^2}(S_t)Z\right| + O(\Delta t^{3/2})$$

Taking conditional expectation and assuming that the expected return from the portfolio is the same as the riskless rate we obtain (we have also removed the order term)

$$\begin{aligned}
r\Delta t\left(-V(s) + s\frac{\partial V}{\partial s}(s)\right) = &-\frac{1}{2}\sigma(s)^2\frac{\partial^2 V}{\partial s^2}(s)\Delta t - g(S_t)\Delta t \\
&- E\left(kS(t)\sigma(S_t)\sqrt{\Delta t}\left|\frac{\partial^2 V}{\partial s^2}(S_t)Z\right|\Big|\mathcal{F}_t\right)
\end{aligned}$$

which is

$$\begin{aligned}
0 = \Bigg(&\frac{1}{2}\sigma(s)^2\frac{\partial^2 V}{\partial s^2}(s) + rs\frac{\partial V}{\partial s}(s) - rV(s) + g(s) \\
&+ \frac{E\left(kS_t\sigma(S_t)\sqrt{\Delta t}|Z|\left|\frac{\partial^2 V}{\partial s^2}(S_t)\right|\,|\mathcal{F}_t\right)}{\Delta t}\Bigg)\Delta t
\end{aligned}$$

The expected value is equal to

$$\int_{-\infty}^{\infty} ks\sigma(s)\sqrt{\Delta t}\left|\frac{\partial^2 V}{\partial s^2}(s)\right||x|\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{x^2}{2}\right)\mathrm{d}x = \sqrt{\frac{2}{\pi}}ks\sigma(s)\sqrt{\Delta t}\left|\frac{\partial^2 V}{\partial s^2}(s)\right|$$

.

We have therefore obtained the non-linear ODE

$$0 = \frac{1}{2}\sigma(s)^2\frac{\partial^2 V}{\partial s^2}(s) + rs\frac{\partial V}{\partial s}(s) - rV(s) + g(s) + \sqrt{\frac{2}{\pi}}ks\sigma(s)\frac{1}{\sqrt{\Delta t}}\left|\frac{\partial^2 V}{\partial s^2}(s)\right| \tag{3.5.14}$$

The argument can also be made when we have bought a claim. The ODE then becomes

$$0 = \frac{1}{2}\sigma(s)^2\frac{\partial^2 V}{\partial s^2}(s) + rs\frac{\partial V}{\partial s}(s) - rV(s) + g(s) - \sqrt{\frac{2}{\pi}}ks\sigma(s)\frac{1}{\sqrt{\Delta t}}\left|\frac{\partial^2 V}{\partial s^2}(s)\right|$$

The model can therefore produce bid and ask prices taking into account delta hedging. For products with a convex or concave value function the increase/decrease in price will be the same as a shift up/down in the volatility. But for a product with a general value function this will not be the case.

Let us work with the ODE (3.5.14). It is on the form (3.4.1) so we discretize it as in (3.4.2). We see that

$$\sigma(s, f''(s))^2 = \frac{1}{2}\sigma(s)^2 + \sqrt{\frac{2}{\pi}}ks\sigma(s)\frac{1}{\sqrt{\Delta t}}\text{sign}\left(f''(s)\right)$$

where we define

$$\text{sign}(x) = \begin{cases} 1 & x \geq 0 \\ -1 & x < 0 \end{cases} \tag{3.5.15}$$

We then have to solve

$$A(F)F = G$$

This we do by the Newton method: Starting with a guess F^1 we can compute a new guess by a linear approximation around in F^1

$$A(F^1)F^1 + D(A(F^1)F^1)(F^2 - F^1) = G \tag{3.5.16}$$

where $D(\cdot)$ denotes the Jacobian, that we need to specify. As $A(F)F$ is not continuously differentiable we will define a generalized Jacobian, see Clarke (1983). We simply specify the derivative as

$$\frac{\partial\sigma(s, f''(s))^2 f''(s)}{\partial f''(s)} = \frac{1}{2}\sigma(s)^2 + \sqrt{\frac{2}{\pi}}ks\sigma(s)\frac{1}{\sqrt{\Delta t}}\text{sign}\left(f''(s)\right)$$

where we use the definition in (3.5.15).

Then we see that

$$D(A(F)F) = A(F)$$

so the left hand side of (3.5.16) is

$$\begin{aligned} A(F^1)F^1 + D(A(F^1)F^1)(F^2 - F^1) &= A(F^1)F^1 + A(F^1)(F^2 - F^1) \\ &= A(F^1)F^2 \end{aligned}$$

The Newton method therefore becomes

- Choose a starting guess F^{old} and set $error = 2\epsilon$.
- While $error > \epsilon$ solve

  $$A(F^{old})F^{new} = G$$

  to obtain F^{new}, and set

  $$\begin{aligned} error &= ||F^{new} - F^{old}|| \\ F^{old} &= F^{new} \end{aligned}$$

Since the function $A(F)F$ is semismooth (it is piecewise linear) it can be shown that the Newton iteration will converge under certain conditions, see Qi and Sun (1993). But parallel to Pooley et al. (2003) we can easily show that the algorithm converge from any starting point if the matrices $A(F^i)$ fulfill the assumptions in proposition 3.3.1:

Consider the two equations

$$\begin{aligned} &A(F^{j-1})F^j = G \Leftrightarrow A(F^j)F^j + (A(F^{j-1}) - A(F^j))F^j = G \\ &A(F^j)F^{j+1} = G \end{aligned}$$

Subtracting them gives us

$$A(F^j)(F^{j+1} - F^j) = (A(F^{j-1}) - A(F^j))F^j$$

Row $0 < i < n$ on the right hand side will have the form

$$\left(\sqrt{\frac{2}{\pi}} k s_i \sigma(s_i) \frac{1}{\sqrt{\Delta t}} \left(\text{sign} \left(\frac{F_{i+1}^{j-1} - 2F_i^{j-1} + F_{i-1}^{j-1}}{\Delta x^2} \right) - \text{sign} \left(\frac{F_{i+1}^{j} - 2F_i^{j} + F_{i-1}^{j}}{\Delta x^2} \right) \right) \right) \left(\frac{F_{i+1}^{j} - 2F_i^{j} + F_{i-1}^{j}}{\Delta x^2} \right) \leq 0$$

as $s_i \geq 0$ and $\sigma(s_i) \geq 0$. Row 0 and n on the right hand side will be 0. Therefore we obtain

$$A(F^j)(F^{j+1} - F^j) \leq 0$$

By proposition 3.3.1 $-A(F^j)$ is an M-matrix which implies $F^{j+1} \geq F^j$. The sequence of solutions is therefore increasing.

Then we show that the sequence is bounded. Each row of

$$A(F^j)F^{j+1} = G$$

can be written as

$$\begin{aligned} &a_i^j F_{i-1}^{j+1} - b_i^j F_i^{j+1} + c_i^j F_{i+1}^{j+1} = G_i \Leftrightarrow \\ &b_i^j F_i^{j+1} = a_i^j F_{i-1}^{j+1} + c_i^j F_{i+1}^{j+1} - G_i \leq a_i^j F_{\max}^{j+1} + c_i^j F_{\max}^{j+1} - G_{\min} \end{aligned}$$

where the subscripts max and min denotes maximal and minimal element of the vector. This implies

$$F_{\max}^{j+1} \leq \frac{-G_{\min}}{|K|}$$

with K defined in proposition 3.3.1.

Then we turn to uniqueness. Assume we have two solutions F^1 and F^2

$$\begin{aligned} &A(F^1)F^1 = G \Leftrightarrow A(F^2)F^1 + \left(A(F^1) - A(F^2)\right)F^1 = G \\ &A(F^2)F^2 = G \end{aligned}$$

subtracting the two equations we obtain

$$A(F^2)\left(F^1 - F^2\right) = (A(F^2) - A(F^1))F^1$$

by the same arguments as above we see that

$$A(F^2)\left(F^1 - F^2\right) \leq 0$$

and as $-A(F^2)$ is an M-matrix we have

$$F^1 \geq F^2$$

interchanging superscripts we get

$$F^1 = F^2.$$

3.5.4.1 Numerical example

We want to look at a perpetual claim with a non convex/concave value function. Therefore let us consider a perpetual range accrual. A range accrual is a product that counts up the number of days a given underlying stays within two boundaries (the range) and pays out an amount in proportion to this number at maturity. Since we consider perpetual claims let us assume the claim pays a continuous stream $q\ \mathrm{d}t$ when the underlying is in the range (a, b). Our differential equation becomes

$$\frac{1}{2}\sigma(s)^2 V''(s) + rsV'(s) - rV(s) = -q1_{\{a<s<b\}}(s) \tag{3.5.17}$$

Including transaction-costs we get

$$\frac{1}{2}\sigma(s)^2V''(s) + rsV'(s) - rV(s) \pm \sqrt{\frac{2}{\pi}}ks\sigma(s)\frac{1}{\sqrt{\Delta t}}\left|V''(s)\right| = -q1_{\{a<s<b\}}(s) \tag{3.5.18}$$

here $\pm$ denotes that we have sold or bought the product. We will assume that a strong comparison principle holds for (3.5.18).

Let us consider a normal Black-Scholes model where $r = 0.05$, $\sigma = 0.2$. Assume $a = 100$, $b = 200$, $k = 0.001$, $q = 0.2$ and that we delta hedge daily, which means $\Delta t = \frac{1}{250}$. In order to price we log-transform the ODE (3.5.18) to obtain

$$\begin{aligned}&\frac{1}{2}\sigma^2\left(\tilde{V}''(x) - \tilde{V}'(x)\right) + r\tilde{V}'(x) - r\tilde{V}(x) \pm \sqrt{\frac{2}{\pi}}k\sigma\frac{1}{\sqrt{\Delta t}}\left|\tilde{V}''(x) - \tilde{V}'(x)\right| \\ &\quad = -q1_{\{\log(a)<x<\log(b)\}}(x)\end{aligned}$$

We discretize the ODE using a uniform grid with 200 grid points between $\log\left(\frac{a+b}{2}\right) - 10\sigma$ and $\log\left(\frac{a+b}{2}\right) + 10\sigma$. We set $\tilde{V}''(x) - \tilde{V}'(x)$ equal to zero on the left boundary and the $\tilde{V}(x)$ equal to zero on the right boundary. This is done in order to make the discrete scheme work.

Note that we have chosen $\frac{k\sigma}{\sqrt{\Delta t}}$ small enough such that (3.4.3) is satisfied.

We will now demonstrate the ability of the model to generate bid/ask prices for non-convex/concave products compared to just shifting the volatility up and down. In figure 3.12 we graph bid, mid and ask prices for a perpetual range accrual found solving (3.5.18). Similarly in figure 3.13 we graph prices obtained solving (3.5.17) where we shift variance up and down by $2\sqrt{\frac{2}{\pi}}k\sigma\frac{1}{\sqrt{\Delta t}}$. We clearly see how nicely the Leland model generates bid and ask prices contrary to just shifting the volatility up and down, this is of course expected.

The model can also be applied to other specifications of the volatility function. We can for example consider a CEV model:

$$\mathrm{d}S_t = rS_t\mathrm{d}t + \sigma S_t^{\beta}\mathrm{d}W_t.$$

This leads to the ODE

$$\frac{1}{2}\sigma^2s^{2\beta}V''(s) + rsV'(s) - rV(s) \pm \sqrt{\frac{2}{\pi}}k\sigma s^{1+\beta}\frac{1}{\sqrt{\Delta t}}\left|V''(s)\right| = -q1_{\{a<s<b\}}(s).$$

Note that monotonicity of the numerical scheme will not be satisfied in this case since the ODE both have a $s^{2\beta}V''(s)$ part and a $s^{1+\beta}\left|V''(s)\right|$ part. From now one we therefore assume the numerical scheme converge to the solution of the ODE.

Let us look at the Bid/Ask spread in the model for different choices of β. We choose

$$\sigma = \tilde{\sigma}\left(\frac{a+b}{2}\right)^{1-\beta}$$

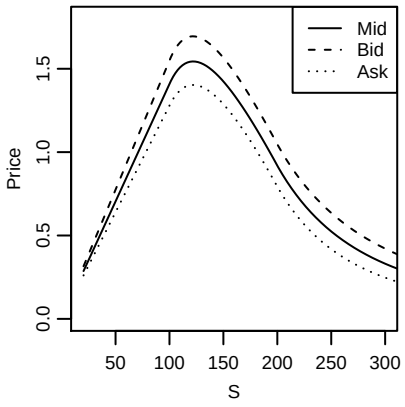

Figure 3.12: Bid, mid and ask prices obtained using the Leland model. Parameters are $r = 0.05$, $\sigma = 0.2$, $a = 100$, $b = 200$, $k = 0.001$, $q = 0.2$ and $\Delta t = \frac{1}{250}$. We use 200 points in a region of -10 to 10 standard deviations.

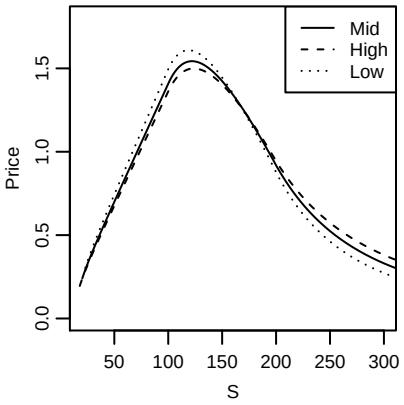

Figure 3.13: Option prices for different levels of volatility. Parameters are $r = 0.05$, $\sigma = 0.2$, $a = 100$, $b = 200$, $k = 0.001$ and $q = 0.2$. We use 200 points in a region of -10 to 10 standard deviations.

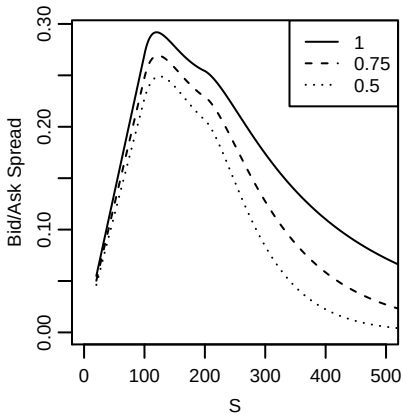

Figure 3.14: Bid/Ask spread for three different CEV models. Parameters are $r = 0.05$, $\sigma = 0.2$, $a = 100$, $b = 200$, $k = 0.001$, $q = 0.2$, $\Delta t = \frac{1}{250}$ and β is varying see the legend. We use 200 points in the region -10 to 10 standard deviations calculated in the log-normal model.

in order to keep everything comparable. Then we set $r = 0.05$, $\sigma = 0.2$, $a = 100$, $b = 200$, $k = 0.001$, $q = 0.2$ and $\Delta t = \frac{1}{250}$.

Figure 3.14 graph the Bid/Ask spread for three different values of β. For models with $\beta \neq 1$ we do not log-transform but use a nonuniform grid constructed as in section 3.5.1.1 with 200 grid points. The endpoints of the grid are specified by -10 and 10 standard deviations from $\frac{a+b}{2}$ in the log-normal model.

As seen a lower CEV coefficient leads to a lower Bid/Ask spread, this is

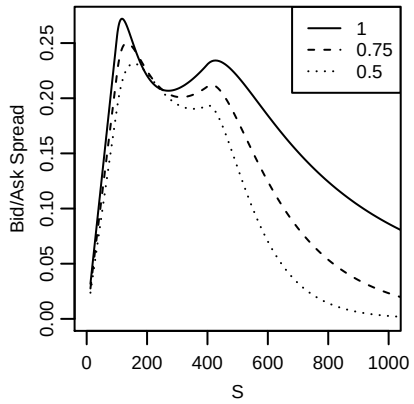

Figure 3.15: Bid/Ask spread for three different CEV models. Parameters are $r = 0.05$, $\sigma = 0.2$, $a = 100$, $b = 400$, $k = 0.001$, $q = 0.2$, $\Delta t = \frac{1}{250}$ and β is varying see the legend. We use 400 points in the region -15 to 15 standard deviations calculated in the log-normal model.

expected for high stock prices since the volatility of the stock will be increasing in β. But for low stock prices the volatility of the stock will be decreasing in β. The lower Bid/Ask spread for lower β is explained by the fact that the absolute gamma in the area around 100 is increasing in β.

If we instead consider a range accrual with a higher right boundary, say $b = 400$ we will see another picture. In figure 3.15 we plot the Bid/Ask spreads for such an option. As seen we now have a range of stock prices where the Bid/Ask spread is decreasing in β.

3.6 Conclusion

This paper has considered a simple finite difference method to solve two-boundary ODEs with financial applications. We prove rigorously that the method converge for some different important cases. On top of this we provide different examples of financial interest where the method is used. Of special interest is the solution to the problem of the value of cash when interest rates can go negative. The numerical solution given here is very simple and we can extend the method to other more complicated models.

Bibliography

Andreasen, J. (2001). Credit explosives. Working paper. Available at SSRN.

Andreasen, J. and D. Bang (2007). Negative interest rates and perpetual cash. Working paper.

Ascher, U. M., R. M. M. Mattheij, and R. D. Russell (1995). *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*. SIAM.

Avellaneda, M., A. Levy, and A. Paras (1995). Princing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance 2*, 73–88.

Barles, G. and P. E. Souganidis (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymtotic Analysis 4*, 271–283.

Berman, A. and R. J. Plemmons (1979). *Nonnegative Matrices In the Mathematical Sciences*. Academic Press.

Björk, T. (2004). *Arbitrage theory in continuous time* (Second ed.). Oxford University Press.

Clarke, F. H. (1983). *Optimization and Nonsmooth Analysis*. John Wiley & Sons, Inc.

Crandall, M. G., H. Ishii, and P.-L. Lions (1992). User's guide to viscosity solutions of second order partial differential equations. *Bullitin of the American Mathematical Society 27*(1), 1–67.

Cryer, C. W. (1983). The efficient solution of linear complementarity problems for tridiagonal minkowski matrices. *ACM Transaction on Mathematical Software 9*(2), 199–214.

D'Halluin, Y., P. A. Forsyth, and G. Labahn (2005). A semi-Lagrangian approach for American Asian options under jump diffusion. *SIAM Journal on Scientific Computing 27*(1), 315–345.

Ekström, E., P. Lötstedt, and J. Tysk (2009). Boundary values and finite difference methods for the single factor term structure equation. *Applied Mathematical Finance 16*(3), 253–259.

Guiso, L., P. Sapienza, and L. Zingales (2013). The determinants of attitudes towards strategic default on mortgages. *Journal of Finance 68*(4), 1473–1515.

Heider, P. (2010). Numerical methods for non-linear black scholes equations. *Applied Mathematical Finance 17*(1), 59–81.

Horn, R. A. and C. R. Johnson (1991). *Topics in Matrix Analysis*. Cambridge University Press.

Karatzas, I. and S. E. Shreve (1998). *Brownian Motion and Stochastic Calculus*. Springer.

Keller, H. B. (1976). *Numerical Solution of Two Point Boundary Value Problems*. Society for Industrial and Applied Mathematics.

Leland, H. E. (1985). Option pricing and replication with transaction costs. *Journal of Finance 40*, 1283–1301.

Lyons, T. (1995). Uncertain volatility and the risk free synthesis of derivatives. *Applied Mathematical Finance 2*, 117–133.

Nabben, R. (1999). Two sided bounds on the inverse of diagonally dominant matrices. *Lin. Alg. Appl. 287*, 289–305.

Pooley, D. M., P. A. Forsyth, and K. R. Vetzal (2003). Numerical convergence properties of option pricing PDEs with uncertain volatility. *IMA journal of Numerical Analysis 23*, 241–267.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (2002). *Numerical recipes in C++* (Second ed.). Cambridge University Press.

Qi, L. and J. Sun (1993). A nonsmooth version of newton's method. *Mathematical Programming 58*, 353–367.

Tavella, D. and C. Randall (2000). *Pricing Financial Instruments, The Finite Difference Method*. John Wiley & Sons, Inc.

Touzi, N. (2010). Optimal stochastics control, stochastic target problems, and backward SDE. Lecture Notes, can be found at http://www.cmap.polytechnique.fr/∼touzi/Fields-LN.pdf.

Wilmott, P., J. Dewynne, and S. Howison (1993). *Option Pricing: Mathematical models and computation*. Oxford Financial Press.

Chapter 4

Fast Ninomiya-Victoir calibration of the Double-Mean-Reverting Model

Christian Bayer[1] Jim Gatheral[2] Morten Karlsmark

Accepted for publication in Quantitative Finance. Available at http://dx.doi.org/10.1080/14697688.2013.818245

Abstract

We consider the three factor double mean reverting (DMR) model of Gatheral (2008), a model which can be successfully calibrated to both VIX options and SPX options simultaneously. One drawback of this model is that calibration may be slow because no closed form solution for European options exists. In this paper, we apply modified versions of the second order Monte Carlo scheme of Ninomiya and Victoir (2008) and compare these to the Euler-Maruyama scheme with partial truncation of Lord et al. (2010), demonstrating on the one hand that fast calibration of the DMR model is practical, and on the other that suitably modified Ninomiya-Victoir schemes are applicable to the simulation of much more complicated time-homogeneous models than may have been thought previously.

4.1 Introduction

It is common knowledge that the Black-Scholes option pricing model is inconsistent with market pricing of options. Local volatility models, Lévy models, stochastic volatility models, stochastic volatility models with jumps and

[1]WIAS Berlin
[2]Department of Mathematics, Baruch College, CUNY

various variants and combinations of these have been proposed to fit market implied volatilities better and describe the dynamics of the resulting volatility surface. With the advent of trading in VIX options in 2006 however, marginal risk-neutral densities of forward volatilities of SPX became effectively observable, substantially constraining possible choices of volatility dynamics. Various authors have since proposed models that price both options on SPX and options on VIX more or less consistently with the market. Notable amongst these are the market models of Bergomi (2005) and the variance curve factor models of Buehler (2006).

In Gatheral (2008), a specific three factor variance curve model was introduced with dynamics motivated by economic intuition for the empirical dynamics of the variance. This model was simultaneously calibrated to SPX and VIX option markets.

In this *double-mean-reverting* or *DMR* model, the dynamics are given by

$$\mathrm{d}S_t = \sqrt{v_t} S_t \mathrm{d}W_t^1, \tag{4.1.1a}$$

$$\mathrm{d}v_t = \kappa_1 \left(v_t' - v_t\right) \mathrm{d}t + \xi_1 \, v_t^{\alpha_1} \, \mathrm{d}W_t^2, \tag{4.1.1b}$$

$$\mathrm{d}v_t' = \kappa_2 \left(\theta - v_t'\right) \mathrm{d}t + \xi_2 \, v_t'^{\alpha_2} \, \mathrm{d}W_t^3, \tag{4.1.1c}$$

where the Brownian motions W_i are all in general correlated with $\mathbb{E}[\mathrm{d}W_t^i \, \mathrm{d}W_t^j] = \rho_{ij} \, \mathrm{d}t$.

Thus variance mean-reverts to a level that itself moves slowly over time with the state of the economy. Also, it is a stylized fact that the distribution of volatility (whether realized or implied) should be roughly lognormal (see Andersen et al. (2001) for example); when the model is calibrated to market option prices, we find that indeed $\alpha_1 \approx 1$ consistent with this stylized fact.

One drawback of this model is that no closed-form solution for European options exists so finite difference or Monte Carlo methods need to be used to price options. Calibration is therefore slow. In Gatheral (2008), the DMR model is calibrated using an Euler-Maruyama Monte Carlo scheme with the partial truncation step of Lord et al. (2010).

In this paper, we show how to apply the second order Monte Carlo scheme of Ninomiya and Victoir (2008) to the calibration of the DMR model, substantially improving calibration time. In passing, we show that a Ninomiya-Victoir second order Monte-Carlo scheme with fully closed-form steps can be achieved for models that are rather more complicated than those (such as the Heston model) to which the technique has been applied so far.

The plan of the paper is as follows. Section 4.2 describes how the model of Gatheral (2008) is calibrated. Section 4.3 explains the Monte Carlo scheme of Ninomiya and Victoir (2008), the drift trick of Bayer et al. (2013) and a subsequent extension by us which we apply to the DMR model. Section 4.4 presents practical examples of calibration to SPX and VIX options with numerical results, and in Section 4.5 we perform a convergence analysis with reasonable parameters. In Section 4.6 we present some concluding remarks.

4.2 Estimating the constants of the DMR model

In Gatheral (2008), the parameters of the DMR model were calibrated to the VIX and SPX options markets with a sequence of steps that we will now individually describe.

4.2.1 Estimation of κ_1, κ_2, θ and ρ_{23}

As of time t, the T-maturity forward variance is given by

$$\xi_t(T) = \mathbb{E}\left[v_T | \mathcal{F}_t\right]$$

and the T-maturity variance swap by

$$\mathbb{E}\left[\int_t^T v_s \, \mathrm{d}s \middle| \mathcal{F}_t\right].$$

Variance swaps are traded in the market so forward variance is a traded asset. Under diffusion assumptions, the fair value of a variance swap is given by evaluating the so-called log-strip of European puts and calls (see Chapter 11 of Gatheral (2006) for example):

$$\mathbb{E}\left[\int_t^T v_s \, \mathrm{d}s \middle| \mathcal{F}_t\right] = 2\left\{\int_{-\infty}^0 p(k)\, dk + \int_0^\infty c(k)\, dk\right\}, \tag{4.2.1}$$

where $k = \log(K/F_{t,T})$ is the log-strike and p and c respectively are put and call prices expressed as a fraction of the strike price. Thus, given a database of historical market option prices, market variance swap prices may be estimated by interpolation, extrapolation and integration.

It is straightforward to verify that in the DMR model (4.1.1), forward variances are given by

$$\xi_t(T) = \theta + (v_t - \theta)\, e^{-\kappa_1 \tau} + (v'_t - \theta)\, \frac{\kappa_1}{\kappa_1 - \kappa_2}\left(e^{-\kappa_2 \tau} - e^{-\kappa_1 \tau}\right), \tag{4.2.2}$$

where $\tau = T - t$. Direct integration then gives us an expression for the spot variance curve[3]

$$\begin{aligned}\mathbb{E}\left[\int_t^T v_s \, \mathrm{d}s \middle| \mathcal{F}_t\right] &= \theta\,\tau + (v_t - \theta)\frac{1 - e^{-\kappa_1 \tau}}{\kappa_1} \\ &\quad + (v'_t - \theta)\frac{\kappa_1}{\kappa_1 - \kappa_2}\left\{\frac{1 - e^{-\kappa_2 \tau}}{\kappa_2} - \frac{1 - e^{-\kappa_1 \tau}}{\kappa_1}\right\}. \end{aligned} \tag{4.2.3}$$

In the DMR model, θ, κ_1 and κ_2 are constants; v_t and v'_t are state variables. From (4.2.3), variance swaps depend linearly on the state variables. Thus, for

[3] (4.2.2) and (4.2.3) may be recognized as the Svensson yield curve model, re-used in our context.

every choice of θ, κ_1 and κ_2, given option prices for various expiries, we can approximate the spot variance curve and infer v_t and v'_t by linear regression.

Performing daily regressions for the seven year period from January 2001 to April 2008, and optimizing over θ, κ_1 and κ_2 to minimize the mean squared error between the fitted curves and actual curves, optimal choices of the parameters θ, κ_1 and κ_2 and also daily time series of v_t and v'_t were obtained. The optimal choice of parameters was found to be

$$\begin{aligned}\theta &= 0.078,\\ \kappa_1 &= 5.5,\\ \kappa_2 &= 0.10.\end{aligned}$$

The correlation ρ_{23} between W_t^2 and W_t^3 is then estimated as the historical correlation between the series v_t and v'_t. The estimated value was

$$\rho_{23} = 0.59.$$

4.2.2 Estimation of the exponents α_1 and α_2

In order to obtain α_1 and α_2 we need information on how the volatility of volatility moves with the volatility itself. To obtain a proxy for the volatility of volatility Gatheral (2008) does the following.

Consider the SABR model for the forward with $\beta = 1$:

$$\begin{aligned}\mathrm{d}F_t &= \alpha_t F_t \mathrm{d}W_t^1,\\ \mathrm{d}\alpha_t &= \nu\alpha_t \mathrm{d}W_t^2,\end{aligned}$$

where $\mathrm{d}W_t^1\mathrm{d}W_t^2 = \rho\mathrm{d}t$. An approximative Black-Scholes volatility for short maturities can be computed using the formula

$$\sigma_{BS}(k) = \alpha_0\, f\left(\frac{k}{\alpha_0}\right) \tag{4.2.4}$$

where $k := \log(K/F_0)$ is the log-strike and

$$f(y) = -\frac{\nu\, y}{\log\left(\frac{\sqrt{\nu^2\, y^2 + 2\,\rho\,\nu\, y + 1} - \nu\, y - \rho}{1-\rho}\right)},$$

see Hagan et al. (2002). It is observed in Gatheral (2008) that the formula can fit observed volatilities very well, even for longer maturities. On a given day we have option quotes for a number of different maturities. We can fit the SABR model using the approximative formula (4.2.4) and obtain coefficients $\alpha_0^\tau, \nu^\tau, \rho^\tau$ for each maturity τ. Gatheral (2008) parametrizes the ν^τ coefficient for the different maturities with the function

$$\nu^\tau = \frac{\nu_{eff}}{\sqrt{\tau}}$$

which fits the term structure of the ν-parameter remarkedly well. The number ν_{eff} is then used as a proxy for the volatility of volatility in a lognormal volatility model on that given day. The VIX index is used as a proxy for volatility.

Using the same dataset as in Section 4.2.1, Gatheral (2008) collects a time-series from January 2001 to April 2008 of ν_{eff} obtained by calibration to SPX options. He also collects VIX quotes. Doing a linear regression of $\log(\nu_{eff})$ onto $\log(\text{VIX})$ he obtains the equation

$$\log(\nu_{eff}) \approx -0.125 - 0.127 \log(\text{VIX}).$$

We can therefore write an SDE for the volatility:

$$\mathrm{d}\alpha_t = c\alpha_t^{-0.127}\alpha_t \mathrm{d}W_t = c\alpha_t^{0.873}\mathrm{d}W_t^2.$$

In the DMR model we are looking for a coefficient on the variance $v_t = \alpha_t^2$. Using Ito's lemma we obtain

$$\mathrm{d}v_t = O(\mathrm{d}t) + 2cv_t^{0.9365}\mathrm{d}W_t.$$

We will use the rounded coefficient $\alpha_1 = 0.94$, which obviously is close to one. There is insufficient market data to be able to say anything about the exponent α_2 so in Gatheral (2008), the choice $\alpha_2 = \alpha_1 = 0.94$ was made. As we will see in Section 4.3.2, various simplifications are possible if $\alpha_1 = \alpha_2 = 1$ (the so-called *double lognormal* model) so that case will also considered in the following.

4.2.3 Daily calibration of remaining parameters

Although the volatility of volatility parameters ξ_1 and ξ_2 are in principle constants of the DMR model, Gatheral (2008) presents empirical evidence that calibrated parameters are not constant in the data. ξ_1 and ξ_2 are thus left free to be calibrated daily to VIX options data. The correlations ρ_{12} and ρ_{13} cannot be imputed from VIX option data; they are left free to improve the daily calibration of the DMR model to SPX data.

So on any given day, both the state variables v_t and v'_t, and the model parameters ξ_1, ξ_2, ρ_{12} and ρ_{13} are calibrated to VIX and SPX options data. v_t and v'_t are calibrated to variance swaps using linear regression and equation (4.2.3). In Gatheral (2008) calibration of ξ_1, ξ_2, ρ_{12} and ρ_{13} was performed using Monte-Carlo simulation. The chosen discretization was an Euler-Maruyama scheme with a partial truncation step, see Lord et al. (2010), which we can write recursively as

$$\begin{aligned}
x((k+1)\Delta) &= -\frac{1}{2}v(k\Delta)\Delta + \sqrt{v(k\Delta)}Z_k^1, \\
\tilde{v}((k+1)\Delta) &= \tilde{v}(k\Delta) + \kappa_2\left(\tilde{v}'(k\Delta) - \tilde{v}(k\Delta)\right)\Delta + \left(\tilde{v}(k\Delta)^+\right)^{\alpha_1} Z_k^2, \\
\tilde{v}'((k+1)\Delta) &= \tilde{v}'(k\Delta) + \kappa_2\left(\theta - \tilde{v}'(k\Delta)\right)\Delta + \left(\tilde{v}'(k\Delta)^+\right)^{\alpha_2} Z_k^3,
\end{aligned}$$

here Δ is the time step, $v(k\Delta) = \tilde{v}(k\Delta)^+$, $v'(k\Delta) = \tilde{v}'(k\Delta)^+$, $x(k\Delta) = \log(S(k\Delta))$, $Z_k^i \sim N(0,\Delta)$ and $\mathbb{E}[Z_k^i Z_k^j] = \rho_{ij}\Delta$. This is a general scheme and we do not need to know moments or asymptotic properties of the density in order to use it. Lord et al. (2010) finds the full truncation scheme superior when simulating the Heston model. When α_1 and α_2 are close to one however, our tests suggest that the partial truncation scheme is superior; this improvement becomes apparent only when time steps are large.

4.3 The Ninomiya-Victoir scheme and drift trick

In Ninomiya and Victoir (2008) a general second order weak discretization scheme for stochastic differential equations was introduced. Consider a multi-dimensional stochastic differential equation in Stratonovich form

$$\mathrm{d}\mathbf{X}(t,\mathbf{x}) = V_0(\mathbf{X}(t,\mathbf{x}))\mathrm{d}t + \sum_{i=1}^{d} V_i(\mathbf{X}(t,\mathbf{x})) \circ \mathrm{d}B_t^i, \tag{4.3.1}$$

where $\mathbf{X}(0,\mathbf{x}) = \mathbf{x} \in \mathbb{R}^N$, $B_t^1, \ldots, B_t^d$ are d *independent* standard Brownian motions and $V_i : \mathbb{R}^N \to \mathbb{R}^N$, $i = 0, \ldots, d$, are sufficiently regular vector fields. In this general setting, the Ninomiya-Victoir scheme based on a uniform grid with time steps Δ is recursively given by

$$\begin{aligned} \mathbf{X}^{(NV)}(0,\mathbf{x}) &= \mathbf{x}, \\ \mathbf{X}^{(NV)}((k+1)\Delta,\mathbf{x}) &= \begin{cases} e^{\frac{\Delta}{2}V_0} e^{Z_k^1 V_1} \cdots e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}(k\Delta,\mathbf{x}), & \Lambda_k = -1, \\ e^{\frac{\Delta}{2}V_0} e^{Z_k^d V_d} \cdots e^{Z_k^1 V_1} e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}(k\Delta,\mathbf{x}), & \Lambda_k = +1. \end{cases} \end{aligned} \tag{4.3.2}$$

Here $e^{tV}\mathbf{x} \in \mathbb{R}^N$ denotes the ODE solution at time $t \in \mathbb{R}$ to

$$\dot{\mathbf{y}} = V(\mathbf{y}), \quad \mathbf{y}(0) = \mathbf{x},$$

i.e., the *flow* of the vector field V,[4] and the probability space carries independent random-variables (Λ_k), with values ± 1 at probability $1/2$, and independent $\mathcal{N}(0,\Delta)$ random variables (Z_k^j). Note that $t = Z_k^j$ can take negative values, so one has to ensure that the ODE solutions used in an implementation of the NV scheme actually do make sense for positive as well as for negative t. One step in the NV scheme corresponds actually to a (non-discrete) cubature formula of order $m = 5$ in the sense of Lyons and Victoir (2004). When seen from this point of view, the reversal of the order of the flows depending on the coin-flip Λ_k serves to improve the approximation of the Lévy area and higher

[4] (4.3.2) is to be read from right to left, i.e., $e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}(k\Delta,\mathbf{x})$ means that solution $e^{\frac{\Delta}{2}V_0} \mathbf{X}^{(NV)}(k\Delta,\mathbf{x})$ of the ODE driven by V_0 is then used as initial value for the ODE driven by the vector field V_d, which is run until the (possibly negative) time Z_k^d.

iterated integrals in the weak sense. On the other hand, one can also interpret the Ninomiya-Victoir scheme as the stochastic version of a classical operator splitting scheme, where the infinitesimal generator $V_0 + \frac{1}{2}\sum_{i=1}^{d} V_i^2$ of the diffusion is split into the first order differential operator V_0 and the second order differential operators $\frac{1}{2}V_1^2, \ldots, \frac{1}{2}V_d^2$.[5] The order-reversal is, in that context, a well known trick which improves the order of the method and goes back to Strang (1963).

The Ninomiya-Victoir scheme has attracted wide attention since its introduction; it is nowadays found in various sophisticated numerical packages such as Inria's software PREMIA for financial option computations. A variation of the scheme designed to deal with degeneracies arising in some affine situations is discussed in Alfonsi (2010).

4.3.1 Improving the efficiency of the Ninomiya-Victoir method

4.3.1.1 Changing the driving noise

In terms of numerical efficiency, cubature methods, and the Ninomiya-Victoir scheme in particular, heavily rely on the ability to solve, fast and accurately, ordinary differential equations. The general cubature methods involve *time-inhomogeneous* ODEs with a rather complicated structure, involving all vector-fields at all times. Thus, there is usually no alternative to solving them numerically, often with Runge-Kutta methods. (A detailed discussion on how Runge-Kutta methods are applied in this context is found in Ninomiya and Ninomiya (2009).)

Using the canonical splitting induced by the model formulation, the Ninomiya-Victoir scheme only involves the composition of solution flows to time-homogeneous ODEs. In particular, there will be "lucky" cases of models where all (or at least most) ODE flows can be solved exactly – in terms of easy-to-evaluate expressions. In such a case, one has effectively found a second order weak approximation method which can be implemented without relying on numerical ODE solvers, and the Ninomiya-Victoir method can be expected to perform especially well in such cases. As was observed by Ninomiya and Victoir (2008), the Heston model is such a lucky case. However, one soon encounters models (e.g., the popular SABR model) in which some of the vector-fields do not allow for flows in closed form. In Bayer et al. (2013), it was found that the class of favorable models can be significantly enlarged by working with an almost trivial modification of the NV scheme. This modification is based on the

[5] Recall that a vector field $V : \mathbb{R}^N \to \mathbb{R}^N$ is identified with the first order linear differential operator acting on smooth functions $f : \mathbb{R}^N \to \mathbb{R}$ by $Vf(x) \equiv \nabla f(x) \cdot V(x)$. By iteration, V^2 can then be interpreted as a linear second order differential operator.

equivalence of (4.3.1) with

$$\begin{aligned}\mathrm{d}\mathbf{X}(t,\mathbf{x}) = &\left(V_0(\mathbf{X}(t,\mathbf{x})) - \sum_{j=1}^{d}\gamma_j V_j(\mathbf{X}(t,\mathbf{x}))\right)\mathrm{d}t + \\ &+\sum_{j=1}^{d} V_j(\mathbf{X}(t,\mathbf{x}))\circ\mathrm{d}\left(B_t^j+\gamma_j t\right) \\ \equiv &V_0^{(\gamma)}(\mathbf{X}(t,\mathbf{x}))\,\mathrm{d}t + \sum_{j=1}^{d} V_j(\mathbf{X}(t,\mathbf{x}))\circ\mathrm{d}\left(B_t^j+\gamma_j t\right)\end{aligned}$$

whatever the choice of drift parameters $\gamma_1,\ldots,\gamma_d$. Assume that all diffusion vector-fields $(V_1,\ldots,V_d)$ allow for flows in closed form, whereas e^{tV_0} is not available in closed form.[6] The point is that, in a variety of concrete examples, one can pick drift parameters $\gamma_1,\ldots,\gamma_d$ in a way that $e^{tV_0^{(\gamma)}}$ can be solved in closed form after all.

Therefore, we propose the following variant of the Ninomiya-Victoir method (which, following Bayer et al. (2013), shall be referred to as the "Ninomiya-Victoir scheme with drift (trick)"):

$$\begin{aligned}&\mathbf{X}^{NVd)}(0,\mathbf{x}) = \mathbf{x}, \\ &\mathbf{X}^{(NVd)}((k+1)\Delta,\mathbf{x}) = \\ &\qquad\begin{cases} e^{\frac{\Delta}{2}V_0^{(\gamma)}}e^{Z_k^1 V_1}\cdots e^{Z_k^d V_d}e^{\frac{\Delta}{2}V_0^{(\gamma)}}\mathbf{X}^{(NVd)}(k\Delta,\mathbf{x}), & \Lambda_k=-1, \\ e^{\frac{\Delta}{2}V_0^{(\gamma)}}e^{Z_k^d V_d}\cdots e^{Z_k^1 V_1}e^{\frac{\Delta}{2}V_0^{(\gamma)}}\mathbf{X}^{(NVd)}(k\Delta,\mathbf{x}), & \Lambda_k=+1,\end{cases}\end{aligned} \tag{4.3.3}$$

where $Z_k^i \sim \mathcal{N}(\Delta\gamma_i,\Delta)$ independent of each other.

Note that (4.3.3) corresponds to the splitting of the differential operator according to

$$V_0 + \frac{1}{2}\sum_{i=1}^{d} V_i^2 = V_0^{(\gamma)} + \sum_{i=1}^{d}\left(\frac{1}{2}V_i^2 + \gamma_i V_i\right).$$

4.3.1.2 Incorporating ODE splitting

The strategy of Section 4.3.1.1, namely to replace the standard Ninomiya-Victoir splitting (4.3.2) by a different one customized to the specific problem at hand, can be generalized to accommodate for an even wider class of problems. In particular, one can directly incorporate any *splitting scheme* (in the ODE sense) for any of the ODEs involved in (4.3.2) into the Ninomiya-Victoir scheme. Let us again assume that (only) the Stratonovich vector field V_0 is

[6] As models are usually devised in the Ito framework and V_0 is obtained from the Ito drift by the Stratonovich correction, this situation is quite common in finance.

too complicated to allow for closed form solutions of the corresponding ODEs. The structure of the Stratonovich drift vector field

$$V_0(\mathbf{x}) = V(\mathbf{x}) - \frac{1}{2}\sum_{i=1}^{d} DV_i(\mathbf{x}) \cdot V_i(\mathbf{x}),$$

where $V : \mathbb{R}^N \to \mathbb{R}^N$ denotes the drift vector field of the SDE in the Ito formulation and DV_i denotes the Jacobian of the vector field V_i, $i = 1, \dots, d$, motivates to apply a classical ODE splitting scheme in order to solve the ODE $\dot{\mathbf{y}} = V_0(\mathbf{y})$, i.e., we try to find vector fields $V_{0,1}$ and $V_{0,2}$ such that $V_0 = V_{0,1} + V_{0,2}$ and the ODEs driven by $V_{0,1}$ and $V_{0,2}$ have (closed-form) solutions $e^{tV_{0,1}}$ and $e^{tV_{0,2}}$, respectively. In that case, the solution $e^{\Delta V_0}$ of the ODE driven by the vector field V_0 at time Δ can be approximated by

$$e^{\Delta V_0}\mathbf{x} = e^{\Delta V_{0,1}} e^{\Delta V_{0,2}}\mathbf{x} + \mathcal{O}(\Delta^2) = e^{\Delta V_{0,2}} e^{\Delta V_{0,1}}\mathbf{x} + \mathcal{O}(\Delta^2),$$

a method sometimes known as *symplectic Euler scheme*, see Hairer et al. (2006). We can incorporate the symplectic Euler method in the Ninomiya-Victoir scheme as follows: starting with $\mathbf{X}^{(NVs)}(0, \mathbf{x}) = \mathbf{x}$, we iterate according to

$$\mathbf{X}^{(NVs)}((k+1)\Delta, \mathbf{x}) = \begin{cases} e^{\frac{\Delta}{2}V_{0,1}} e^{\frac{\Delta}{2}V_{0,2}} e^{Z_k^1 V_1} \cdots e^{Z_k^d V_d} e^{\frac{\Delta}{2}V_{0,2}} e^{\frac{\Delta}{2}V_{0,1}} \mathbf{X}^{(NVs)}(k\Delta, \mathbf{x}), & \Lambda_k = -1, \\ e^{\frac{\Delta}{2}V_{0,1}} e^{\frac{\Delta}{2}V_{0,2}} e^{Z_k^d V_d} \cdots e^{Z_k^1 V_1} e^{\frac{\Delta}{2}V_{0,2}} e^{\frac{\Delta}{2}V_{0,1}} \mathbf{X}^{(NVs)}(k\Delta, \mathbf{x}), & \Lambda_k = +1. \end{cases} \tag{4.3.4}$$

Even though the symplectic Euler scheme only has local order two, the Strang trick of repeating the symplectic Euler scheme once while inverting the order of the vector fields, again produces a scheme with local order three and, hence, global order two. Indeed, note that the *Verlet scheme*

$$e^{\Delta V_0}\mathbf{x} = e^{\frac{\Delta}{2}V_{0,1}} e^{\Delta V_{0,2}} e^{\frac{\Delta}{2}V_{0,1}}\mathbf{x} + \mathcal{O}(\Delta^3)$$

obtained by omitting the diffusion part in (4.3.4) has (global) order two. Both the Verlet and the symplectic Euler scheme are examples of *geometric integrators* for ODEs, and we again refer to Hairer et al. (2006) for much more information.

4.3.1.3 Analysis of the modified Ninomiya Victoir scheme of (4.3.4)

We give a sketch of the proof that the modified NV algorithm (4.3.4) has second order convergence in the weak sense. In particular, in the following we assume sufficient regularity for all involved functions and vector fields. Let $\mathbf{X}_\Delta$ denote the true solution of the SDE at time Δ and let $\overline{\mathbf{X}}_\Delta \equiv \mathbf{X}^{(NVs)}(\Delta, \mathbf{x})$

denote the output of the modified NV algorithm after one time step of size Δ, both started at $\mathbf{x}$ at time 0. By the Markov property, it suffices to show that the weak local error is of third order (see, for instance, Talay and Tubaro (1990)), i.e., that

$$\mathbb{E}[f(\mathbf{X}_\Delta)] - \mathbb{E}[f(\overline{\mathbf{X}}_\Delta)] = \mathcal{O}(\Delta^3) \tag{4.3.5}$$

for sufficiently smooth test functions f. Indeed, denote
$u(t,\mathbf{y}) = \mathbb{E}\left[g\left(\mathbf{X}_T\right) \mid \mathbf{X}_t = \mathbf{y}\right]$ and assume that we want to approximate $\mathbb{E}\left[g(\mathbf{X}_T)\right] = u(0,\mathbf{x})$ by $\mathbb{E}\left[g(\overline{\mathbf{X}}_{n\Delta})\right] = \mathbb{E}\left[u(n\Delta, \overline{\mathbf{X}}_{n\Delta})\right]$ with $\Delta = T/n$. Then, by a telescoping sum, we may decompose the global error $\mathbb{E}\left[g(\overline{\mathbf{X}}_{n\Delta})\right] - u(0,\mathbf{x})$ as the sum of the local errors

$$\mathbb{E}\left[g(\overline{\mathbf{X}}_{n\Delta})\right] - u(0,\mathbf{x}) = \sum_{k=0}^{n-1} \mathbb{E}\left[u((k+1)\Delta, \overline{\mathbf{X}}_{(k+1)\Delta}) - u(k\Delta, \overline{\mathbf{X}}_{k\Delta})\right].$$

Assuming (4.3.5), we have

$$\mathbb{E}\left[u((k+1)\Delta, \overline{\mathbf{X}}_{(k+1)\Delta}) - u(k\Delta, \overline{\mathbf{X}}_{k\Delta})\right] = \mathcal{O}(\Delta^3),$$

by first conditioning on $\overline{\mathbf{X}}_{k\Delta}$. However, we sum $n = T/\Delta$ of these terms, so that the global error is $\mathcal{O}(\Delta^2)$.

For the proof of (4.3.5), note that the "multiplication" of vector fields in the sense of iterative applications of vector fields as differential operators is certainly non-commutative. Taylor expansion applied to (4.3.4) implies that

$$\begin{aligned}
\mathbb{E}[f(\overline{\mathbf{X}}_\Delta)] = \frac{1}{2} &\left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2\right)\left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2\right) \\
&\left(1 + \frac{1}{2}\Delta V_1^2 + \frac{1}{8}\Delta^2 V_1^4\right) \cdots \left(1 + \frac{1}{2}\Delta V_d^2 + \frac{1}{8}\Delta^2 V_d^4\right) \\
&\left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2\right)\left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2\right) f(\mathbf{x}) + \\
&+ \frac{1}{2}\left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2\right)\left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2\right) \\
&\left(1 + \frac{1}{2}\Delta V_d^2 + \frac{1}{8}\Delta^2 V_d^4\right) \cdots \left(1 + \frac{1}{2}\Delta V_1^2 + \frac{1}{8}\Delta^2 V_1^4\right) \\
&\left(1 + \frac{1}{2}\Delta V_{0,2} + \frac{1}{8}\Delta^2 V_{0,2}^2\right)\left(1 + \frac{1}{2}\Delta V_{0,1} + \frac{1}{8}\Delta^2 V_{0,1}^2\right) f(\mathbf{x}) + \mathcal{O}(\Delta^3) \\
= f(\mathbf{x}) &+ \Delta\left(V_0 + \frac{1}{2}\sum_{i=1}^{d} V_i^2\right) f(\mathbf{x}) + \\
&+ \frac{1}{2}\Delta^2\left(V_0 + \frac{1}{2}\sum_{i=1}^{d} V_i^2\right)^2 f(\mathbf{x}) + \mathcal{O}(\Delta^3) \\
= \mathbb{E}[f(\mathbf{X}_\Delta)] &+ \mathcal{O}(\Delta^3),
\end{aligned}$$

where the last equality follows since $V_0 + \frac{1}{2}\sum_{i=1}^{d} V_i^2$ is the infinitesimal generator of the diffusion $\mathbf{X}(t, \mathbf{x})$.

4.3.2 The Ninomiya-Victoir scheme for the DMR model

4.3.2.1 The Stratonovich formulation of the DMR model

Consider again the DMR model (4.1.1) re-expressed in terms of independent Brownian motions B^i:

$$\begin{aligned}
\mathrm{d}S_t &= \sqrt{v_t}\, S_t\, \mathrm{d}B_t^1, \\
\mathrm{d}v_t &= \kappa_1 \left(v_t' - v_t\right) \mathrm{d}t + \xi_1 v_t^{\alpha_1} \left(\tilde{\rho}_{1,2} \mathrm{d}B_t^1 + \sqrt{1 - \tilde{\rho}_{1,2}^2}\, \mathrm{d}B_t^2\right), \\
\mathrm{d}v_t' &= \kappa_2 \left(\theta - v_t'\right) \mathrm{d}t + \xi_2 v_t'^{\alpha_2} \left(\tilde{\rho}_{1,3} \mathrm{d}B_t^1 + \tilde{\rho}_{2,3} \mathrm{d}B_t^2 + \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2}\, \mathrm{d}B_t^3\right),
\end{aligned} \tag{4.3.6}$$

where $\tilde{\rho}_{12} = \rho_{12}$, $\tilde{\rho}_{13} = \rho_{13}$ and $\tilde{\rho}_{23} = \frac{\rho_{23} - \rho_{12}\rho_{13}}{\sqrt{1-\rho_{12}^2}}$. In order to apply the simulation method of Ninomiya and Victoir (2008) we need to re-express the Itô SDEs (4.3.6) in Stratonovich form (see for example Definition 3.13 of Karatzas and Shreve (1988)).

We first compute the quadratic covariation terms as follows:

$$\begin{aligned}
d\left[\sqrt{v_t}\, S, B^1\right]_t &= \left\{\frac{1}{2}\, \tilde{\rho}_{1,2}\, v_t^{\alpha_1 - \frac{1}{2}} + v_t\right\} S_t\, \mathrm{d}t \\
d\left[\xi_1 v_t^{\alpha_1}, \left(\tilde{\rho}_{1,2}\, B^1 + \sqrt{1 - \tilde{\rho}_{1,2}^2}\, B^2\right)\right]_t &= \xi_1^2\, \alpha_1\, v_t^{2\alpha_1 - 1}\, \mathrm{d}t \\
d\left[\xi_2 v_t'^{\alpha_2}, \left(\tilde{\rho}_{1,3}\, B^1 + \tilde{\rho}_{2,3}\, B^2 + \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2}\, B^3\right)\right]_t &= \xi_2^2\, \alpha_2\, v_t'^{2\alpha_2 - 1}\, \mathrm{d}t.
\end{aligned}$$

We then obtain the Stratonovich form of (4.3.6):

$$\mathbf{X}(t, \mathbf{x}) = \mathbf{x} + \int_0^t V_0\left(\mathbf{X}(s, \mathbf{x})\right)\, ds + \sum_{j=1}^{3} \int_0^t V_j\left(\mathbf{X}(s, \mathbf{x})\right) \circ dB_s^j \tag{4.3.7}$$

where the state vector $\mathbf{X}(t, \mathbf{x}) = (S_t, v_t, v_t')^T$, the initial condition $\mathbf{x} = (S_0, v_0, v_0')^T$, and the driving vector fields are given by

$$V_0(\mathbf{x}) = \begin{pmatrix} -\frac{1}{2}\left(\frac{1}{2}\, \xi_1\, \tilde{\rho}_{1,2}\, x_2^{\alpha_1 - \frac{1}{2}}\, x_1 + x_2\, x_1\right) \\ -\kappa_1\left(x_2 - x_3\right) - \frac{1}{2}\, \xi_1^2\, \alpha_1\, x_2^{2\alpha_1 - 1} \\ -\kappa_2(x_3 - \theta) - \frac{1}{2}\, \xi_2^2\, \alpha_2\, x_3^{2\alpha_2 - 1} \end{pmatrix},$$

and

$$\begin{aligned}
V_1(\mathbf{x}) &= \begin{pmatrix} \sqrt{x_2}\, x_1 & \tilde{\rho}_{1,2}\, \xi_1\, x_2^{\alpha_1} & \tilde{\rho}_{1,3}\, \xi_2\, x_3^{\alpha_2} \end{pmatrix}^T, \\
V_2(\mathbf{x}) &= \begin{pmatrix} 0 & \sqrt{1 - \tilde{\rho}_{1,2}^2}\, \xi_1\, x_2^{\alpha_1} & \tilde{\rho}_{2,3}\, \xi_2\, x_3^{\alpha_2} \end{pmatrix}^T, \\
V_3(\mathbf{x}) &= \begin{pmatrix} 0 & 0 & \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2}\, \xi_2\, x_3^{\alpha_2} \end{pmatrix}^T.
\end{aligned}$$

In order to proceed with the Ninomiya-Victoir splitting, we thus need to solve the ordinary differential equations

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{x}(t) = V_i(\mathbf{x}(t))$$

for all $i \in \{0, 1, 2, 3\}$ and $t \in \mathbb{R}$ with some given boundary condition.

4.3.2.2 The flow of the diffusion vector fields

Following Bayer et al. (2013), it is straightforward to verify that the solution to the (one-dimensional) ODE

$$\frac{\mathrm{d}}{\mathrm{d}t}x(t) = h(t)^\alpha\, x(t)^\beta \tag{4.3.8}$$

is given by

$$x(t) = \begin{cases} \left[(1-\beta)H(t) + x(0)^{1-\beta}\right]_+^{\frac{1}{1-\beta}}, & 0 < \beta < 1, \\ x(0)\, e^{H(t)}, & \beta = 1, \end{cases} \tag{4.3.9}$$

where

$$H(t) = \int_0^t h(s)^\alpha \,\mathrm{d}s.$$

We can thus solve the ODEs

$$\frac{\mathrm{d}}{\mathrm{d}t}x(t) = V_i(x(t))$$

for $i \in \{1, 2, 3\}$ in closed form. The NV algorithm requires solutions of the ODEs driven by the diffusion vector fields for negative times t. As $\frac{\mathrm{d}}{\mathrm{d}t}x(-t) = -\dot{x}(-t)$, this essentially means that the sign of the coefficient changes. In any case, for a fixed time interval $I = [0, T]$ (or $I = [-T, 0]$ in the negative time case), (4.3.9) is the unique solution to (4.3.8), provided that $x(0) \neq 0$. This follows by standard arguments when $x > 0$ on I. On the other hand, note that in our case $h(t)$ cannot change its sign on I for the ODEs under considerations here. Thus, H is always a monotonous function on I. So, for $x(0) > 0$, we can only get $x(t) = 0$ for some $t \in I$, if H is negative on I. But then x must stay at 0 for the remaining time to T (or $-T$, respectively). If, however, $x(0) = 0$, then there might, indeed, be several real-valued solutions for (4.3.8). For instance in the positive time case, when $h > 0$, both (4.3.9) and $x \equiv 0$ are solutions. This phenomenon reflects the fact that the NV scheme should not be expected to perform any better than the EM scheme, when paths of the underlying SDE come too close to 0 too often, as has been found in several numerical studies, e.g., by Lord et al. (2010).

4.3.2.3 The flow of the Stratonovich drift vector field

Solving the ODE for $i = 0$ is a little trickier with no obvious closed-form solution for general values of the exponents α_i. Whereas a numerical solution would be possible by for example applying a Runga-Kutta method, by further splitting the operator, we may obtain a reasonably simple closed-form simulation step. That is, we write

$$V_0 = V_{0,1} + V_{0,2}$$

with

$$V_{0,1}(x) = \begin{pmatrix} -\frac{1}{2}\, x_2\, x_1 \\ -\kappa_1\,(x_2 - x_3) \\ -\kappa_2\,(x_3 - \theta) \end{pmatrix}, \quad V_{0,2}(x) \quad = \begin{pmatrix} -\frac{1}{4}\,\xi_1\,\tilde{\rho}_{1,2}\, x_2^{\alpha_1 - \frac{1}{2}}\, x_1 \\ -\frac{1}{2}\,\xi_1^2\,\alpha_1\, x_2^{2\alpha_1 - 1} \\ -\frac{1}{2}\,\xi_2^2\,\alpha_2\, x_3^{2\alpha_2 - 1} \end{pmatrix}.$$

and solve the equations

$$\frac{\mathrm{d}}{\mathrm{d}t} x(t) = V_{0,j}(x(t)) \text{ with } j = 1, 2.$$

Solution for $j = 1$

The equation in the third row which reads

$$\frac{\mathrm{d}}{\mathrm{d}t} x_3(t) = -\kappa_2\,(x_3 - \theta)$$

has the solution

$$x_3(t) = \theta + e^{-\kappa_2\, t}\,(x_3(0) - \theta)\,.$$

The second ODE has the solution

$$x_2(t) \quad = \quad e^{-\kappa_1\, t}\, x_2(0) + \kappa_1 \int_0^t e^{-\kappa_1\,(t-s)}\, x_3(s)\,\mathrm{d}s$$

which is just the forward variance curve $\xi_t(T)$. The first ODE reads

$$\frac{\mathrm{d}}{\mathrm{d}t} x_1(t) = -\frac{1}{2}\, x_2(t)\, x_1(t)$$

with the solution

$$x_1(t) = x_1(0)\, \exp\left\{-\frac{1}{2}\, H(t)\right\}$$

with $H(t) = \int_0^t x_2(s)\,\mathrm{d}s$ which we recognize as the spot variance swap curve.

Solution for $j = 2$

If $\alpha_1 \neq 1$ and $\alpha_2 \neq 1$, the second and third ODEs have solutions

$$\begin{aligned} x_2(t) &= \left[x_2(0)^{2(1-\alpha_1)} - \alpha_1 (1-\alpha_1) \xi_1^2 t\right]_+^{\frac{1}{2(1-\alpha_1)}}, \\ x_3(t) &= \left[x_3(0)^{2(1-\alpha_2)} - \alpha_2 (1-\alpha_2) \xi_2^2 t\right]_+^{\frac{1}{2(1-\alpha_2)}}. \end{aligned}$$

The first ODE reads

$$\frac{\mathrm{d}}{\mathrm{d}t} x_1(t) = -\frac{1}{4} \xi_1 \tilde{\rho}_{1,2} x_2^{\alpha_1 - \frac{1}{2}} x_1$$

with the solution

$$x_1(t) = x_1(0) \exp\left\{-\frac{1}{4} \xi_1 \tilde{\rho}_{1,2} H(t)\right\}$$

where $H(t) = \int_0^t x_2(s)^{\alpha_1 - \frac{1}{2}} \mathrm{d}s$. $H(t)$ can also then be computed explicitly as

$$H(t) = \frac{2}{\alpha_1 (3/2 - \alpha_1) \xi_1^2} \left\{x_2(0)^{3/2-\alpha_1} - x_2(t)^{3/2-\alpha_1}\right\}.$$

4.3.2.4 The double lognormal case: $\alpha_1 = \alpha_2 = 1$

The special case $\alpha_1 = 1$, $\alpha_2 = 1$ gives the double lognormal model of Gatheral (2008), a model which both fits the empirical SPX and VIX surfaces well and displays remarkable parameter stability. In this case, we may obtain even simpler closed-form solutions by applying the drift trick of Bayer et al. (2013) explained in Section 4.3. This trick involves simplifying V_0 by subtracting components spanned by the the vector fields V_1, V_2, V_3. This is achieved by introducing drift in the Brownian motions. Specifically, with

$$V_0^\gamma = V_0 - \gamma_1 V_1 - \gamma_2 V_2 - \gamma_3 V_3,$$

and choosing

$$\begin{aligned} \gamma_1 &= -\xi_1 \tilde{\rho}_{1,2}, \\ \gamma_2 &= -\frac{\kappa_1 + \frac{1}{2}\xi_1^2 + \gamma_1 \tilde{\rho}_{1,2} \xi_1}{\xi_1 \sqrt{1 - \tilde{\rho}_{1,2}^2}}, \\ \gamma_3 &= -\frac{\kappa_2 + \frac{1}{2}\xi_2^2 - \tilde{\rho}_{1,3} \xi_2 \gamma_1 - \tilde{\rho}_{2,3} \xi_2 \gamma_2}{\xi_2 \sqrt{1 - \tilde{\rho}_{1,3}^2 - \tilde{\rho}_{2,3}^2}}, \end{aligned}$$

we have the much simpler expression

$$V_0^\gamma = \begin{pmatrix} -\frac{1}{2} x_2 x_1 \\ \kappa_1 x_3 \\ \kappa_2 \theta \end{pmatrix}.$$

The third and second differential equations respectively have solutions

$$x_3(t) = x_3(0) + \kappa_2\,\theta\,t,$$
$$x_2(t) = x_2(0) + \kappa_1\,\left(x_3(0)\,t + \frac{1}{2}\kappa_2\,\theta\,t^2\right).$$

The first ODE reads

$$\frac{\mathrm{d}}{\mathrm{d}t}x_1(t) = -\frac{1}{2}\,x_2(t)\,x_1(t)$$

with the solution

$$x_1(t) = x_1(0)\exp\left\{-\frac{1}{2}H(t)\right\}$$

where

$$\begin{aligned} H(t) &= \int_0^t x_2(s)\,\mathrm{d}s \\ &= x_2(0)\,t + \kappa_1\,\left(\frac{1}{2}\,x_3(0)\,t + \frac{1}{6}\kappa_2\,\theta\,t^3\right). \end{aligned}$$

4.3.3 Summary of the modified Ninomiya-Victoir procedure

Adopting the notation of Bayer et al. (2013), if Λ_k is a Bernoulli random variable, the kth NV time step of length Δ in the modified NV simulation of Section 4.3.2.3 is of the form

$$\begin{aligned} &\mathbf{X}\left((k+1)\,\Delta, \mathbf{x}\right) \\ = &\begin{cases} e^{\frac{1}{2}\Delta V_{0,1}}\,e^{\frac{1}{2}\Delta V_{0,2}}\,e^{Z_k^1 V_1}\,e^{Z_k^2 V_2}\,e^{Z_k^3 V_3}\,e^{\frac{1}{2}\Delta V_{0,2}}\,e^{\frac{1}{2}\Delta V_{0,1}}\mathbf{X}(k\,\Delta, \mathbf{x}) & \text{if } \Lambda_k = -1, \\ e^{\frac{1}{2}\Delta V_{0,1}}\,e^{\frac{1}{2}\Delta V_{0,2}}\,e^{Z_k^3 V_3}\,e^{Z_k^2 V_2}\,e^{Z_k^1 V_1}\,e^{\frac{1}{2}\Delta V_{0,2}}\,e^{\frac{1}{2}\Delta V_{0,1}}\mathbf{X}(k\,\Delta, \mathbf{x}) & \text{if } \Lambda_k = +1 \end{cases} \end{aligned}$$

where the $Z_k^i \sim N(0, \Delta)$ are independent of each other.

Similarly, in the NV procedure with drift trick of Section 4.3.2.4, the kth NV time step is of the simpler form

$$\begin{aligned} &\mathbf{X}\left((k+1)\,\Delta, \mathbf{x}\right) \\ = &\begin{cases} e^{\frac{1}{2}\Delta V_0^\gamma}\,e^{Z_k^1 V_1}\,e^{Z_k^2 V_2}\,e^{Z_k^3 V_3}\,e^{\frac{1}{2}\Delta V_0^\gamma}\mathbf{X}(k\,\Delta, \mathbf{x}) & \text{if } \Lambda_k = -1 \\ e^{\frac{1}{2}\Delta V_0^\gamma}\,e^{Z_k^3 V_3}\,e^{Z_k^2 V_2}\,e^{Z_k^1 V_1}\,e^{\frac{1}{2}\Delta V_0^\gamma}\mathbf{X}(k\,\Delta, \mathbf{x}) & \text{if } \Lambda_k = +1 \end{cases} \end{aligned}$$

where the $Z_k^i \sim N(\gamma_i\,\Delta, \Delta)$ (note the nonzero drift) are independent of each other.

4.4 Calibrating the model

4.4.1 Calibrating daily parameters

As mentioned in Section 4.2.3 we would like to infer v_t, v_t', ξ_1, ξ_2, ρ_{12} and ρ_{13} daily. The calibration of these parameters is divided into several steps.

4.4.1.1 v_0 and v_0'

We saw in Section 4.2.1 that the prices of variance swaps may be estimated from the market prices of SPX options using equation (4.2.1). This computation requires a continuous set of option prices which we obtain by fitting the SVI parametrization (see for example Gatheral and Jacquier (2012)) to the volatility smile for each expiry. Then from (4.2.3), given κ_1, κ_2 and θ, v_0 and v_0' may be obtained by linear regression.

4.4.1.2 ξ_1 and ξ_2

The volatility parameters ξ_1 and ξ_2 of the variance processes are obtained by calibrating the model to the market prices of VIX options.

We proxy the underlying of a VIX option by the expected forward variance in our model. The payoff of a call option on the VIX index with strike K expiring at time T may therefore be written as

$$\left(\sqrt{\mathbb{E}\left[\int_T^{T+\Delta} v_s \mathrm{d}s \,\middle|\, \mathcal{F}_T \right]} - K \right)^+$$

where Δ is the length (approximately one month) of the VIX index. For each Monte Carlo path we have a value for v_T and v_T', so the expected forward variance $\mathbb{E}\left[\int_T^{T+\Delta} v_s \mathrm{d}s \,\middle|\, \mathcal{F}_T \right]$ may be computed using (4.2.3). Averaging over all paths gives the model price of the VIX option.

Our chosen objective function is the sum of squared difference between market VIX option prices and the model VIX option prices, both expressed in terms of Black-Scholes implied volatility. Errors are weighted by the reciprocal of the bid-ask spread:

$$\sqrt{\sum_i \left(\frac{\sigma_i^{mid} - \sigma_i^{model}}{\sigma_i^{ask} - \sigma_i^{bid}} \right)^2}.$$

The minimization is performed with a Levenberg-Marquardt algorithm, setting starting values to $\xi_1 = 2.5$ and $\xi_2 = 0.4$, values typical of those that we observe.

4.4.1.3 ρ_{12} and ρ_{13}

We are then left with the two parameters ρ_{12} and ρ_{13} to calibrate. These are used to fit the SPX volatility surface. Our chosen objective function is again the sum of squared differences between market SPX option prices and model SPX option prices, all in implied volatility terms and weighted by the reciprocal of the bid-ask spread.

The objective function can though have multiple local minima, which lead to poor performance of the Levenberg-Marquardt algorithm. This is especially true when using the EM algorithm in combination with pseudo random numbers. To improve performance we need to find a good starting point before applying the solver. We achieve this by evaluating the function at a number of points and starting the Levenberg-Marquardt algorithm at the best point.

When we fit the Heston model to SPX option data, the imputed stock-volatility correlation parameter is typically around -0.7. It seems reasonable then that the two correlation parameters ρ_{12} and ρ_{13} should be in the same ballpark. We therefore search for ρ_{12} and ρ_{13} in the region $[-1,-0.5]\times[-1,-0.5]$, restricted by condition that $\tilde{\rho}_{13}^2+\tilde{\rho}_{23}^2\leq 1$. A good starting point can be found by evaluating the objective function at 30 Sobol points in this region.

4.4.2 Calibration examples with tests of Monte Carlo schemes

We now consider two calibration examples: One with data from April 3, 2007, and the other with data from September 15, 2011. In both of these examples, we will compare the calibration performance of the modified Ninomiya-Victoir scheme described in Section 4.3 with that of the Euler-Maruyama scheme with partial truncation step described in Section 4.2.3. We thereby test both the model and the calibration routines in pre- and post-crisis environments.

Our testing strategy is as follows: With respectively $2^9=512, 2^{11}=2048$ and $2^{13}=8192$ paths we calibrate the model using $6,10,20,\ldots,100,200,300,400,500,1000,2000$ time steps, thereby obtaining calibrated values for ξ_1 and ξ_2.

For each such calibration we obtain optimal volatility parameters. For example $\xi_1^{opt,30,11}$ and $\xi_1^{opt,30,11}$, are optimal parameters for a calibration with 30 time steps and 2^{11} paths. Using the optimal parameters we then perform another Ninomiya-Victoir Monte Carlo simulation with $2^{16}=65,536$ paths and 500 time steps. This latter simulation we use to measure how well the parameters obtained by the calibration fit the market. This is done by calculating the mean squared error objective function (RMSE) between the model prices obtained using $2^{16}=65,536$ paths, 500 time steps and the market prices. We can then assess what the minimum required number of paths and time steps is to obtain a calibration accurate enough for practical applications.

We also compare the performance of a classic Monte Carlo (MC) scheme using pseudo random numbers with that of a Quasi Monte Carlo (QMC) scheme using Sobol quasi random numbers. If the dimension of the integration problem is small, Quasi Monte Carlo should theoretically result in a lower integration error compared to Monte Carlo. Increasing the number of dimensions however decreases the efficiency of the Quasi Monte Carlo method and at some point the Monte Carlo method beats it with a lower integration error. In our implementation, the dimension of the integration problem is the number of time

steps times the number of random variables required per time step (which is three for EM and four for NV due to the coin flip).

There exists a number of heuristic ways to deal with the curse of dimensionality of QMC, see for example da Silva and Barbe (2005). These methods will though not be tried out.

All computations were run on an ASUS desktop with an Intel Core i3 cpu at 2.40 GHz CPU and 4GB memory. The simulations were done in Java using the SSJ package, see "http://www.iro.umontreal.ca/ simardr/ssj/indexe.html". For the Sobol sequences we used the built in direction numbers up to 360 dimensions. Sequences with more dimension were created using direction numbers from the webpage "http://web.maths.unsw.edu.au/ fkuo/sobol/new-joe-kuo-6.21201", these have been obtained using the search criteria $D^{(6)}$ see Joe and Kou (2008). Pseudo random numbers were generated using the Mersenne twister, MT19937. Optimization were done using the Levenberg-Marquardt algorithm present in the SSJ package, this is a Java translation of the MINPACK routine, see More et al. (1980).

4.4.3 April 3, 2007

4.4.3.1 The data

The SPX option dataset contains prices for 421 options, 388 of them include both bid and ask prices, we only use these options in our calibration. There are 14 different option maturities in the dataset ranging from 0.005 to 2.71 years, the forward for the first maturity is 1438.62 and for the last maturity 1556.75. The strikes for the different options lie in the interval 600 to 2000. The VIX option dataset contains prices for 108 options, 96 of them include both bid and ask prices, again we only use options with both bid and ask prices. The dataset contains 7 different maturities ranging from 0.04 years to 1.13 years. The forward for the first maturity is 13.97 and 15.29 for the last maturity. Strikes lie in the interval 10 to 30.

4.4.3.2 Calibration of $\mathbf{v_0}$ and $\mathbf{v_0'}$

As explained in Section 4.4.1.1, we use linear regression to calibrate model variance swaps to market variance swaps (proxied by the SVI log-strip) giving us the parameters

$$v_0 = 0.0153,$$
$$v_0' = 0.0224.$$

The resulting fit is graphed in Figure 4.1.

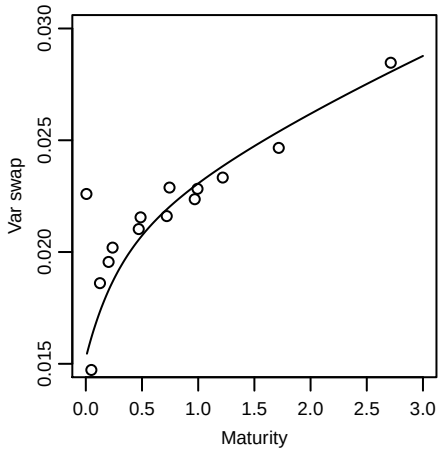

Figure 4.1: SPX market variance swaps as points together with the calibrated model variance swap curve (solid line). Data from April 3, 2007.

4.4.3.3 Calibration of ξ_1 and ξ_2 to VIX options

In Figure 4.2 we have graphed the RMSE from the different calibrations as a function of number of time steps used.

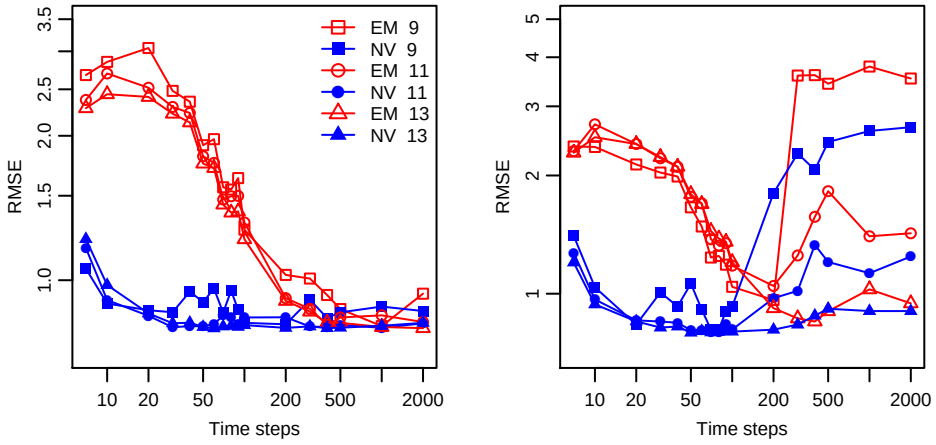

(a) Pseudo random numbers (b) Quasi random numbers

Figure 4.2: RMSE from NV and EM calibration of the DMR model to VIX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (a) specifies the method and the $\log_2$ number of paths.

The figure clearly indicate that a lower bound for the calibration RMSE exists around 0.8. Even with 2^{13} paths we do not get below this barrier. We see a clear advantage of using the NV scheme compared to the EM scheme. Using the NV scheme a calibration can be done using $20 - 30$ time steps and

2^{11} or 2^{13} paths. In comparison we need $400 - 500$ time steps if we use the EM scheme. Therefore we can reduce the number of time steps by a factor of 15 or so.

The integration error seems to be negligible at 2^{11} paths. A good compromise between calibration quality and computational cost therefore seems to be an NV scheme using 2^{11} paths and 30 time steps.

In Figure 4.3 we graph market VIX Black-Scholes implied volatility smiles together with model smiles. The model parameters ξ_1 and ξ_2 were calibrated to the market using MC with 30 NV time steps and 2^{11} paths. Total calibration time was 1.47 seconds. The resulting calibrated ξ parameters are:

$$\xi_1 = 2.873,$$
$$\xi_2 = 0.302.$$

Model option prices were then computed using 100 NV QMC time steps and 2^{16} paths. By inspection of Figure 4.3, the quality of the calibration is quite

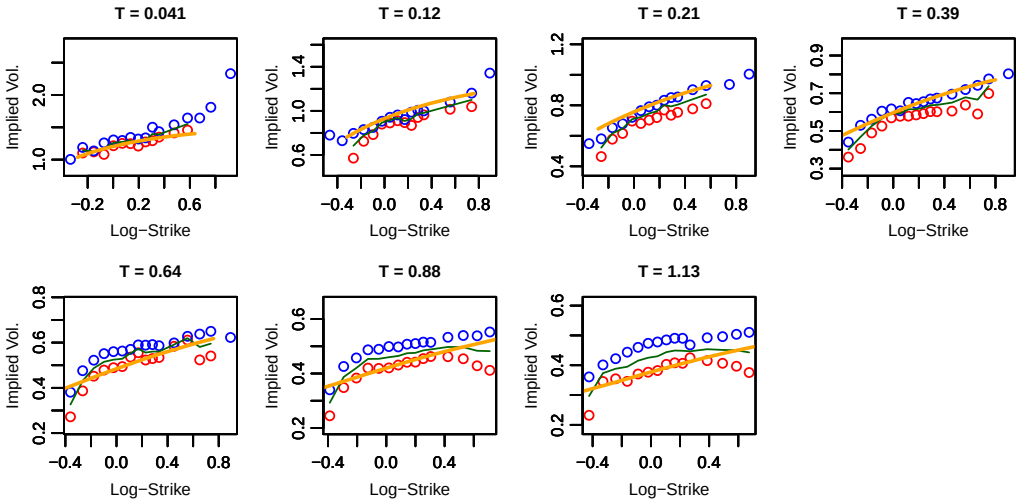

Figure 4.3: Implied Black volatilities for VIX options on April 3, 2007 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from a QMC-NV scheme using 100 time steps and 2^{16} paths (yellow line). The model parameters ξ_1 and ξ_2 are obtained by a calibration using a MC-NV scheme with 30 time steps and 2^{11} Monte Carlo paths.

acceptable, though VIX option bid-ask spreads are admittedly wide.

4.4.3.4 Calibration of ρ_{12} and ρ_{13} to SPX options

This test consists in fixing all parameters including the values of ξ_1 and ξ_2 found by the MC-NV calibration of section 4.4.3.3.

Let us start by doing a normal Levenberg-Marquadt calibration started in $(-0.7, -0.7)$, without an initial search for an optimal point. Figure 4.4 shows the RMSE results

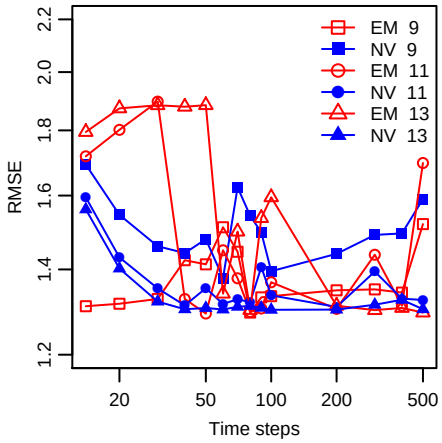

Figure 4.4: RMSE from NV and EM calibration of the DMR model to SPX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$, $\xi_1 = 2.873$, $\xi_2 = 0.302$ and the calibration is done with quasi random numbers using a Levenberg-Marquardt optimizer starting in $(-0.7, -0.7)$. The legend specifies the method and the $\log_2$ number of paths.

The NV scheme performs acceptably, especially using 2^{13} paths, but the 2^9 paths calibration seems to be unacceptably off. The EM scheme shows strange behavior: note that for the first number of time steps the 2^9 paths calibration does best of all the methods, while the EM scheme with more paths shows jumpy behavior, sometimes it finds a good minimum other times it does not. In order to improve the method we use the simple search algorithm described in Section 4.4.1.3. In Figure 4.5 we have applied this before the Levenberg-Marquardt optimizer. After the search algorithm has been applied the EM scheme performs just as well if not better than the NV scheme. The NV scheme with 2^{13} paths and 30 time steps could in principle save us from using the search algorithm but since the NV scheme is much more involved than the EM scheme we will not obtain a speedup. The RMSEs reported in Figure 4.5 clearly show that 2^{11} QMC paths are sufficient in order to obtain a good calibration. The spikes in figure 4.5(a) seems strange, but they only exists for a low number of paths and for pseudo random number. We conclude that it is best to use quasi random numbers when calibrating to SPX options.

In Figure 4.6, we graph market SPX Black-Scholes implied volatility smiles together with model smiles. The correlation parameters ρ_{12} and ρ_{13} were calibrated to the market using QMC with 30 EM time steps and 2^{11} paths. Total calibration time was 2.94 seconds in this case. The resulting calibrated

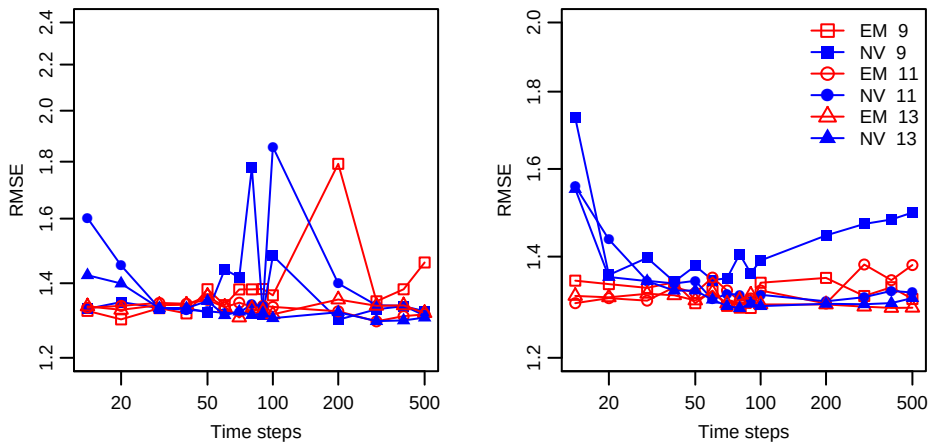

(a) Pseudo random numbers (b) Quasi random numbers

Figure 4.5: RMSE from NV and EM calibration of the DMR model to SPX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$, $\xi_1 = 2.873$, $\xi_2 = 0.302$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (b) specifies the method and the $\log_2$ number of paths.

parameters are:

$$\rho_{12} = -0.992,$$
$$\rho_{13} = -0.615$$

Model SPX option prices were again computed using a QMC-NV scheme with 100 time steps and 2^{16} paths.

Inspecting Figure 4.6, we note that the DMR model fits longer expiration SPX option smiles very well, shorter expirations somewhat less well. This of course is not unexpected for a three-factor model of the DMR type. Note the kinks in the model implied volatilities for $T = 0.049$ and 0.20. These appear since we use too few paths, the model will therefore produce a time value of zero for options far from ATM. Finally, we observe that the calibrated parameters are quite consistent with those reported in Gatheral (2008), except ρ_{12}. However the calibration routine described here runs very much faster.

4.4.3.5 Calibration RMSE

One could argue that the preceeding tests only examine whether or not the methods produce parameters that are good, and not that the schemes with the chosen number of paths and time steps fit the market prices. In figure 4.7 we graph the RMSE between the market prices and the model prices calculated with the same number of paths and time steps that we use in the calibrations. We have for example obtained $\xi_1^{opt,30,11}$ and $\xi_1^{opt,30,11}$ the optimal parameters for a calibration with 30 time steps and 2^{11} paths. To obtain the calibration

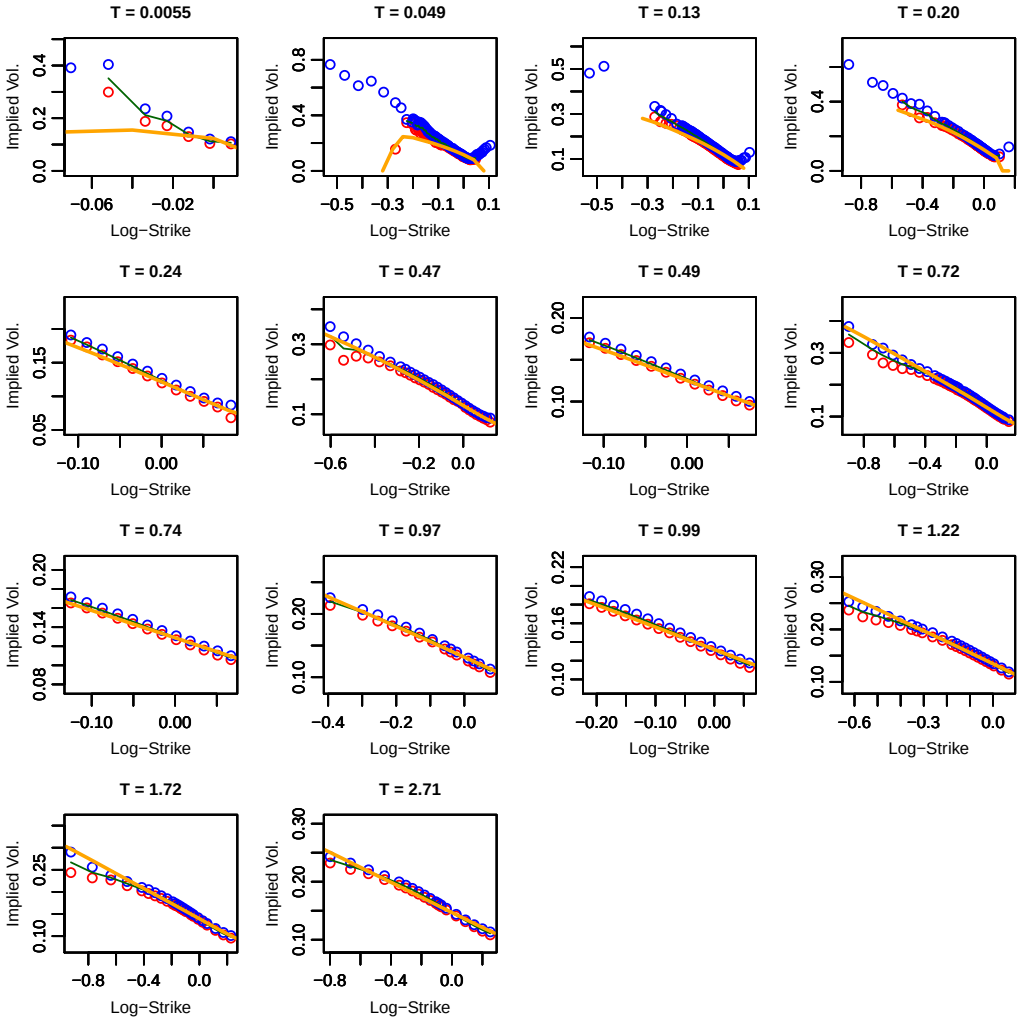

Figure 4.6: Implied Black volatilities for SPX options on April 3, 2007 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from a QMC-NV scheme using 100 time steps and 2^{16} paths (yellow line). The parameters for the model have been obtained using an NV scheme with 30 time steps and 2^{11} paths to calibrate the model to the VIX options, and an EM scheme with 30 time steps and 2^{11} paths to calibrate the model to the SPX options.

RMSE we first simulate the model using 30 time steps, 2^{11} paths, $\xi_1^{opt,30,11}$ and $\xi_1^{opt,30,11}$. Then we calculate the RMSE between these model prices and market prices.

For the VIX options we do not observe that big a difference from the previous RMSE graph. The RMSE only become a bit more bumpy, which is expected since we have higher integration error. The NV scheme with 2^{13} or 2^{11} paths can fit market prices with relatively few steps while the EM scheme

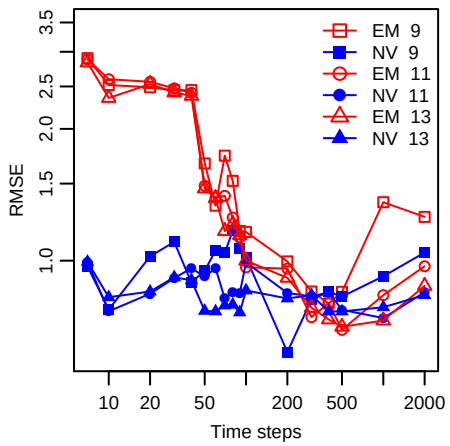

(a) Vix options, pseudo random numbers

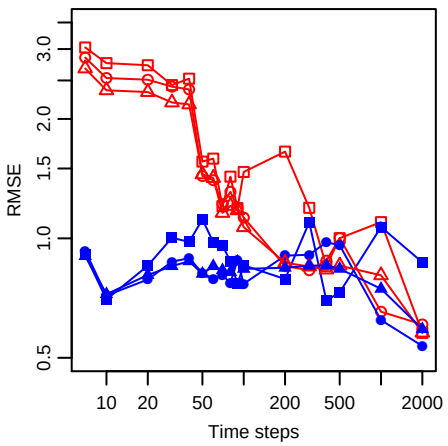

(b) Vix options, quasi random numbers

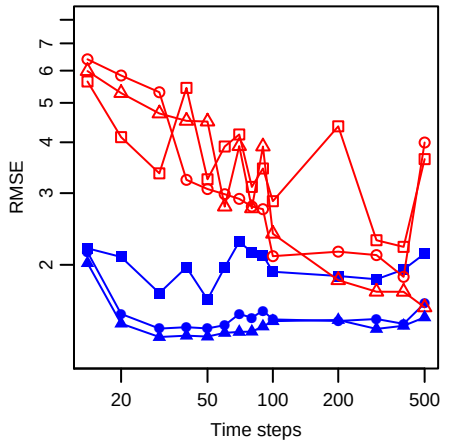

(c) SPX options, quasi random numbers but no search algorithm

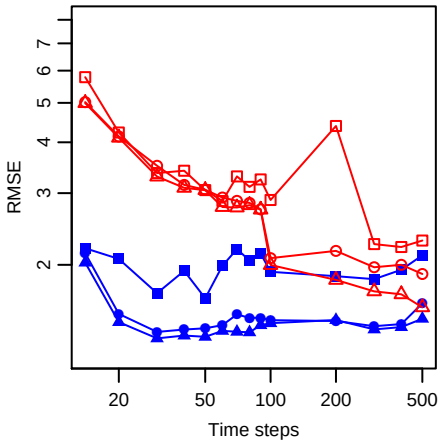

(d) SPX options, quasi random numbers and search algorithm

Figure 4.7: RMSE from NV and EM calibration of the DMR model to VIX and SPX option prices from the 3rd of April 2007. We have set $\alpha_1 = \alpha_2 = 0.94$ and for the calibration to SPX options we have also set $\xi_1 = 2.873$, $\xi_2 = 0.302$. Calibration to VIX options are considered in (a) with pseudo random numbers and in (b) with quasi random numbers. Calibration to SPX options are considered in (c) with quasi random numbers but without the search algorithm and in (d) with quasi random number and the search algorithm. The legend in (a) specifies the method and the $\log_2$ number of paths.

needs at least 200 steps.

For SPX options we observe that the NV scheme with 2^{13} or 2^{11} paths and relatively few time steps generates good fits to market prices. In contrast, the EM scheme requires a large number of time steps to achieve the same fit quality. Nevertheless, if we are concerned only with calibration, the EM scheme does produce perfectly acceptable parameters with few time steps at

significantly lower computational cost than NV.

4.4.4 September 15, 2011

The dataset of Section 4.4.3 is from a period before the 2008 financial crisis. To investigate whether the model still works under more recent market conditions, and to further compare simulation schemes, we now calibrate the DMR model to data from September 15, 2011.

4.4.4.1 The data

The SPX option dataset contains prices for 1176 options, 931 of them with both bid and ask prices, which we use in our calibration. The dataset contains 14 different maturities ranging from 0.0027 to 2.26 years. The strikes range from 100 to 3000. The forward for the first maturity is 1207.70 and for the last maturity 1159.83.

The VIX option data contains prices for 202 options, 148 of them have both bid and ask prices, again these are options we consider in our calibration. The dataset contains 6 different maturities ranging from 0.016 years to 0.42 years. The strikes ranges from 17 to 100. The forward is 33.23 for the first maturity and 31.29 for the last maturity.

4.4.4.2 Calibration of the parameters

Using linear regression to calibrate v_0 and v_0' gives us the parameters

$$v_0 = 0.114,$$
$$v_0' = 0.110.$$

We then test how well the various simulation schemes calibrate the model to VIX options as in Section 4.4.3.3, presenting the results in figure 4.8.

The NV discretization again beats the EM discretization but the difference is not as significant as in the 2007 example, requiring of the order of a factor 10 fewer time steps. The reasons for the decrease in relative performance are twofold: There are no long-dated VIX options in the 2011 dataset, and the volatility processes have higher starting values.

The longest VIX options have a maturity of 1.13 years in the 2007 dataset and 0.42 years in the 2011 dataset. We therefore expect the EM scheme to use $\frac{1.13}{0.42} = 2.69$ as many time steps in 2007 than in 2011, everything else being equal. Since the NV-scheme is second order convergent it will only use $\sqrt{2.69} = 1.64$ as many time steps in 2007 than in 2011.

The higher value for the volatility processes means that they will hit zero less frequently. This affects both schemes in a positive way. We can therefore reduce the number of time steps for both schemes compared to the 2007 calibration. But there exists a lower bound, the number of VIX option maturities

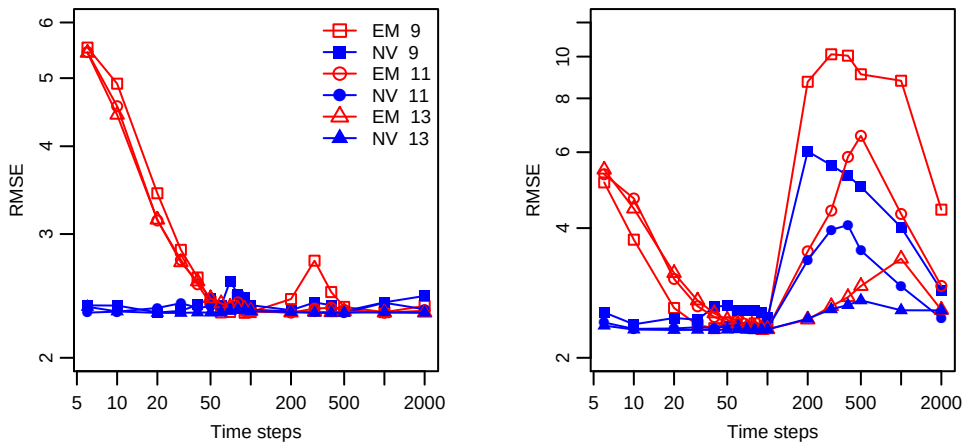

(a) Pseudo random numbers (b) Quasi random numbers

Figure 4.8: RMSE from NV and EM calibration of the DMR model to VIX option prices from the 15th of September 2011. We have set $\alpha_1 = \alpha_2 = 0.94$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (a) specifies the method and the $\log_2$ number of paths.

we need to hit. Therefore the number of time steps cannot be lower than 6, even though the NV scheme could approximate the longest maturity well with fewer equidistant time steps.

In Figure 4.9 we graph market VIX Black-Scholes implied volatility smiles together with model smiles. The calibrated ξ parameters are

$$\xi_1 = 2.689,$$
$$\xi_2 = 0.502.$$

they were obtained using an MC-NV scheme with 6 time steps and 2^{11} paths. Model option prices were then computed using 100 NV QMC time steps and 2^{16} paths.

We observe that the DMR model generates VIX smiles that are too flat. This suggests that the lognormal DMR model with $\alpha_1 = \alpha_2 = 1$, which is faster to simulate, may also fit better. We investigate this in Section 4.4.5. The bid-ask spread have also decreased compared to the 2007 data, making it harder for the model to hit the market prices.

Using ξ_1 and ξ_2 just obtained we then calibrate correlation parameters. Results from this calibration are presented in Figure 4.10.

Again, we see no advantage in using the NV discretization over the simpler (and less costly) EM discretization and we also conclude that quasi random numbers have to be used when calibrating to SPX options.

In Figure 4.11 we graph market SPX Black-Scholes implied volatility smiles together with model smiles. The model parameters ρ_{12} and ρ_{13} were calibrated

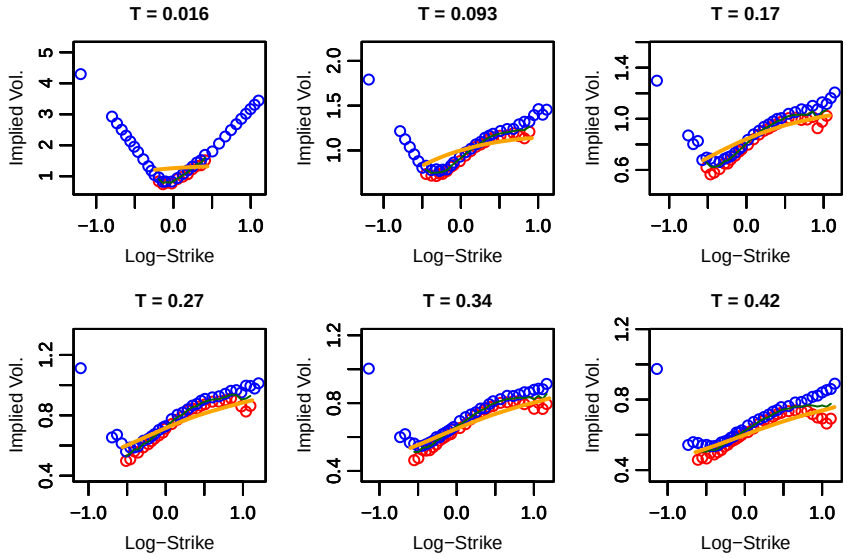

Figure 4.9: Implied Black volatilities for VIX options on September 15, 2011 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from a QMC-NV scheme using 100 time steps and 2^{16} paths (yellow line). The model parameters ξ_1 and ξ_2 are obtained by a calibration using a NV scheme with 6 time steps and 2^{11} Monte Carlo paths.

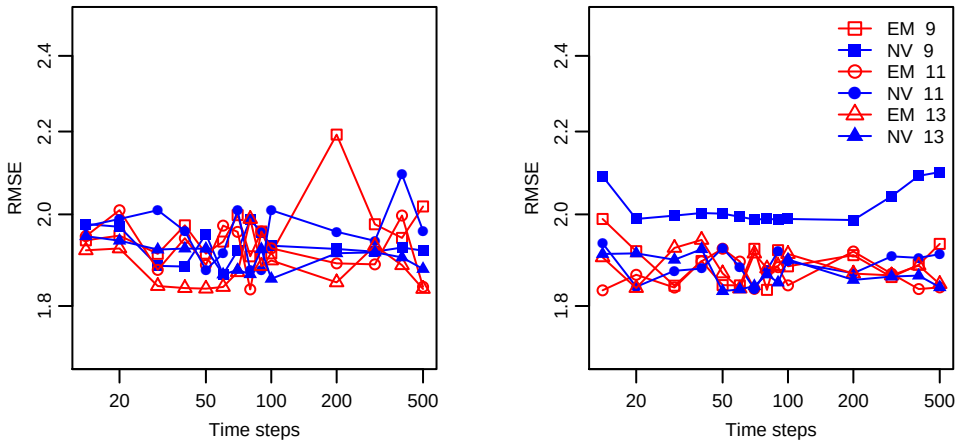

(a) Pseudo random numbers (b) Quasi random numbers

Figure 4.10: RMSE from NV and EM calibration of the DMR model to SPX option prices from the 15th of September 2011. We have set $\alpha_1 = \alpha_2 = 0.94$, $\xi_1 = 2.689$ and $\xi_2 = 0.502$. Pseudo random number are used in (a) and quasi random numbers are used in (b). The legend in (b) specifies the method and the $\log_2$ number of paths.

to the market using QMC-EM scheme with 14 time steps, 2^{11} paths and search

for a good starting point. The resulting calibrated ρ parameters are:

$$\rho_{12} = -0.982,$$
$$\rho_{13} = -0.727.$$

Model option prices were then computed using 100 NV QMC time steps and 2^{16} paths. As with the 2007 calibration, the DMR model fits SPX option prices well except for very short expirations.

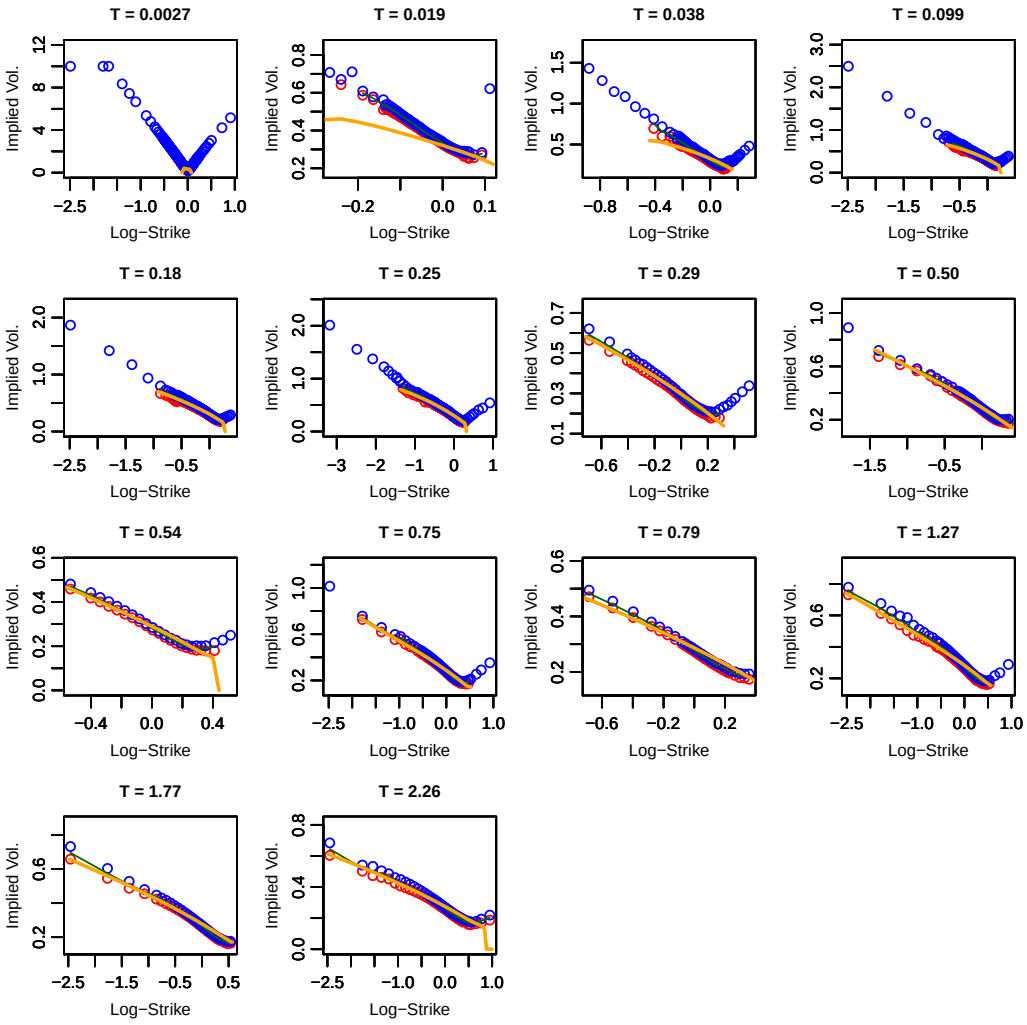

Figure 4.11: Implied Black volatilities for SPX options on September 15, 2011 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from the NV scheme using 100 time steps and 2^{16} Quasi Monte Carlo paths (yellow line). The parameters for the model have been obtained using the NV scheme with 6 time steps and 2^{11} Monte Carlo paths to calibrate the model to the VIX options, and an EM scheme with 14 time steps and 2^{11} Quasi Monte Carlo paths to calibrate the model to the SPX options.

4.4.5 Lognormal DMR model calibration to 2011 data

In the previous section we saw that a model where $\alpha_1 = \alpha_2 = 0.94$ generates VIX option smiles that are too flat compared to the market prices of our 2011 example. In order to increase the steepness of the smile we calibrate the simpler lognormal DMR model with $\alpha_1 = \alpha_2 = 1$.

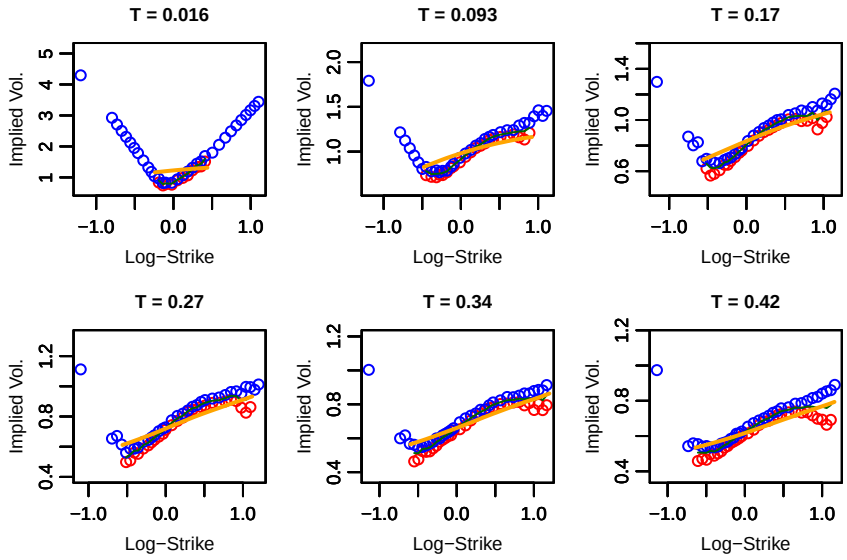

Figure 4.12: Implied Black volatilities for VIX option prices on September 15, 2011 (bid price (red dots), ask price (blue dots) and mid price (green line)) and model prices from QMC-NV with 100 time steps and 2^{16} paths (yellow line). The parameters are obtained by calibration using 10 NV time steps and 2^{11} QMC paths.

In Figure 4.12 we graph the VIX option smiles obtained from the lognormal model. The volatility parameters have been obtained using an MC-NV scheme with 10 time steps and 2^{11} paths. The smiles have steepened compared to the graph in Section 4.4.4.2, the lognormal DMR model therefore fits the market better than the more complicated DMR model calibrated in Section 4.4.4.2. But the option smiles still seem to be too flat. Calibrating the double lognormal model to the SPX options yields more or less the same smiles as before.

4.4.6 Computation times

As we have seen, the NV-scheme can reduce by a large factor the number of time steps needed to achieve a good calibration of the DMR model to VIX options. In the 2007 example we got a reduction of a factor 15 and in the 2011 example we got a reduction of a factor 10. But this will not lead to an equivalent reduction in computation time because the NV discretization involves more computation and is therefore slower.

In Table 4.1, we present the empirical computational cost of the NV discretization relative to the EM discretization. The results are obtained using quasi random numbers, we get almost the same results if we use pseudo random numbers. Because of the drift trick, the NV discretization step is much simpler in the case $\alpha_1 = \alpha_2 = 1$ therefore we present this case separately. 2D is the case where we only simulate the variance processes i.e. when we have to price VIX options. 3D is the simulation of the full model.

	2D	3D
$\alpha_1 = \alpha_2 = 0.94$	4.55	6.84
$\alpha_1 = \alpha_2 = 1$	1.81	3.08

Table 4.1: Relative computation times for NV steps in terms of EM steps. 2D means simulation of the variance process only (*i.e.* for VIX options); 3D means simulation of the full model. The values are obtained by simulating with 90 time steps and 2^{18} QMC paths using the parameters obtained in the 2011 calibrations.

We conclude that it is better to use the EM discretization when calibrating to SPX options where there is little if any RMSE reduction benefit from using the NV step. However, for VIX options, we can achieve a speedup of $3-4$ times in the 2007 example, 2 in the 2011 example and 5 in the 2011 lognormal DMR example. In summary, the optimal calibration recipe appears to be:

- Calibrate ξ_1 and ξ_2 with a Ninomiya-Victoir scheme.
- Calibrate ρ_{12} and ρ_{13} with an Euler-Maruyama scheme.

Using Java code with 30 time steps and 2^{11} paths we can typically calibrate the model to both SPX and VIX option markets in approximately 5 seconds.

4.5 Convergence of the discretization schemes

Having demonstrated in Section 4.4 that we have fast and accurate calibration of the DMR model to VIX and SPX options and moreover that fits to the market are good, we focus in this section on numerical tests of the convergence of the Ninomiya-Victoir (NV) discretization scheme presented in Section 4.3.2 relative to that of the Euler-Maruyama scheme with partial truncation presented in Section 4.2.3.

We use parameters resulting from the calibrations of Section 4.4.4 summarized in Table 4.2. We consider options with a maturity of one year and three different strikes, 0.8 times the forward, ATM and 1.2 times the forward. All option prices are computed using randomized QMC. The randomization is done by scrambling the net and then adding a random shift to the QMC numbers, see Glasserman (2004). This is done in order to obtain a Monte Carlo error around the price.

θ	0.078
κ_1	5.5
κ_2	0.1
ρ_{23}	0.59
v_0	0.114
v_0'	0.110
α_1	0.94
α_2	0.94
ξ_1	2.689
ξ_2	0.502
$\rho_{12} = \tilde{\rho}_{12}$	-0.982
$\rho_{13} = \tilde{\rho}_{13}$	-0.727
$\tilde{\rho}_{23}$	-0.656

Table 4.2: Parameters from the calibration to data from September 15, 2011 with $\alpha_1 = \alpha_2 = 0.94$.

The "true" option prices are computed using 8 independent realizations, each realization is calculated using the NV scheme with 200 time steps and $2^{27} = 134,217,728$ QMC paths.

To obtain convergence graphs we calculate option prices using $5, 10, 20, 30, 50, 70$ time steps for the NV scheme and $5, 10, 20, 30, 50, 70, 100, 200, 500$ time steps for the EM scheme. We simulate 64 independent realizations of the NV prices and 128 independent realizations of the EM prices, each price is computed using $2^{23} = 8.388.608$ QMC paths. Convergence graphs can be seen in Figure 4.13.

The VIX option graphs in Figure 4.13 clearly show that the NV scheme has second order convergence while the EM scheme only converges with order one. For SPX options the picture is blurred a bit by the strange behavior of the EM scheme. There seems to be a kink in the error graph around 50 time steps. The kink exists because the EM scheme with a small number of time steps creates too high option prices, while the EM scheme with a large number of steps creates too low option prices. Therefore the EM scheme price has to cross the "true" price at one point, this point lies around 50 time steps. After the kink we see first order convergence of the EM scheme. The NV scheme clearly shows second order convergence.

In Lord et al. (2010) the Ninomiya-Victoir scheme was found inferior to the full truncation scheme when simulating the Heston model. In the Heston model the stochastic volatility is a CIR process ($\alpha_1 = 0.5$ and no v_t'), but in the DMR model we consider $\alpha_1 \approx 1$. If the parameters of the CIR process violate the Feller condition, 0 is an attainable boundary, and even if they satisfy the Feller condition the process can hit 0 when we simulate the model discretely. If

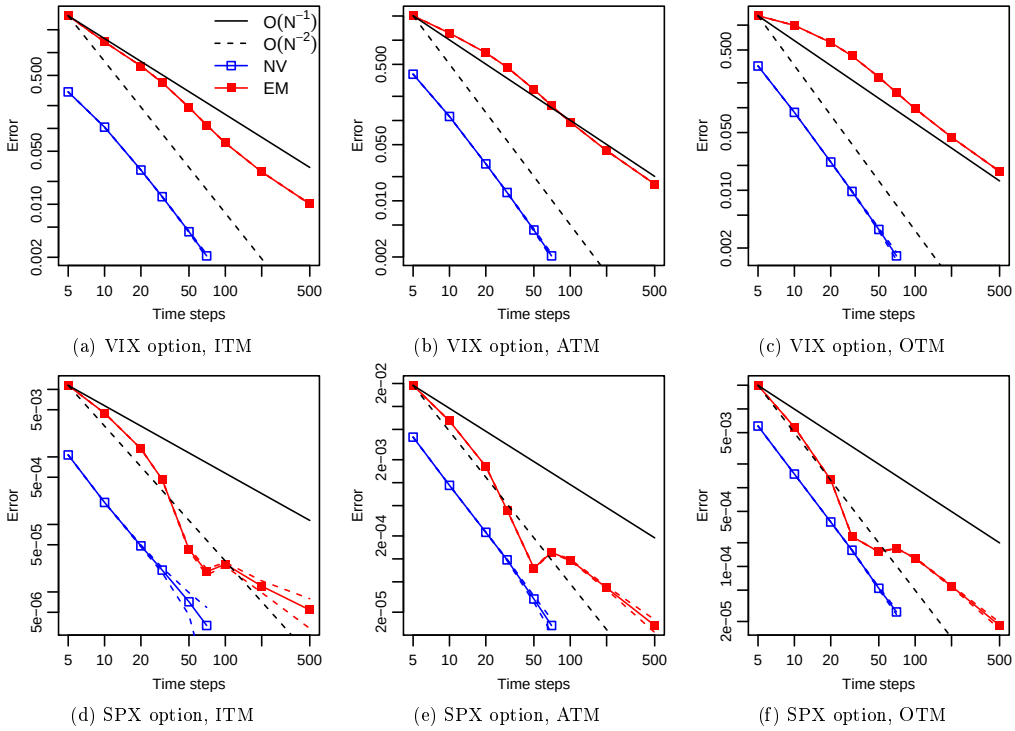

(a) VIX option, ITM (b) VIX option, ATM (c) VIX option, OTM

(d) SPX option, ITM (e) SPX option, ATM (f) SPX option, OTM

Figure 4.13: Pricing error for 1 year VIX and SPX options in a DMR model with parameters from 15th of September 2011 where $\alpha_1 = \alpha_2 = 0.94$. A confidence interval of two standard deviations around the error are marked with the dashed lines. The legend in (a) specifies the method.

α_1 is closer to one this will happen less frequently. We may therefore expect a method which is second order accurate for smooth volatility and drift functions with bounded derivatives to perform better when $\alpha_1 \approx 1$.

Figure 4.13 (d)-(f) clearly show that the Ninomiya-Victoir scheme comes much closer to the true price of SPX options for a given number of time steps, this is also what we conclude from section 4.4.3.5. But the extra precision of the NV scheme will be of no use when calibrating the correlation parameters since we can do that using the EM scheme with a low number of time steps. Given its significantly lower computational cost, the EM scheme is therefore to be preferred when calibrating to SPX options.

As for pricing VIX options, (which Lord et al. (2010) do not consider), our tests show that the outperformance of the NV scheme is sufficient for it to be preferred over the EM scheme for the calibration of volatility parameters.

4.6 Conclusion

In this paper, we have presented two straightforward modifications of the standard Ninomiya-Victoir discretization scheme that conserve second order weak convergence but permit simple closed-form solutions to the ODE's, avoiding the use of numerical integration methods such as Runge-Kutta. Using these schemes for VIX options and the simpler Euler-Maruyama scheme for SPX options, we demonstrated that it is possible to achieve fast and accurate calibration of the DMR model of Gatheral (2008) to both SPX and VIX options markets simultaneously. Moreover, we demonstrated that the DMR model fits SPX and VIX options market data well for two particular dates chosen to represent two very different market environments from before and after the 2008 financial crisis. The fitted parameters of the model over time appear to be remarkably stable.

Finally we performed an empirical analysis of the convergence of the Ninomiya-Victoir (NV) and Euler-Maruyama discretization schemes demonstrating that the NV scheme was indeed second-order weak convergent.

Acknowledgments

We are very grateful to the anonymous referees for their valuable and helpful comments.

Bibliography

Alfonsi, A. (2010). High order discretization schemes for the CIR process: application to affine term structure and Heston models. *Math. Comp. 79*(269), 209–237.

Andersen, T., T. Bollerslev, F. Diebold, and H. Ebens (2001). The distribution of realized stock return volatility. *Journal of Financial Economics 61*(1), 43–76.

Bayer, C., P. Friz, and R. Loeffen (2013). Semi-closed form cubature and applications to financial diffusion models. *Quantitative Finance 13*(5), 769–782.

Bergomi, L. (2005, October). Smile dynamics II.

Buehler, H. (2006). Consistent variance curve models. *Finance and Stochastics 10*, 178–203.

da Silva, M. E. and T. Barbe (2005). Quasi-Monte Carlo in finance: Extending for problems of high effective dimension. *Economia Aplicada 9*(4).

Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. John Wiley & Sons, Inc.

Gatheral, J. (2008). Consistent modelling of SPX and VIX options. Bachelier congress, London.

Gatheral, J. and A. Jacquier (2012). Arbitrage-free SVI volatility surfaces. SSRN.

Glasserman, P. (2004). *Monte Carlo Methods In Financial Engineering*. Springer.

Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002, September). Managing smile risk. *Wilmott Magazine 1*(1), 84–108.

Hairer, E., C. Lubich, and G. Wanner (2006). *Geometric numerical integration* (Second ed.), Volume 31 of *Springer Series in Computational Mathematics*. Springer-Verlag. Structure-preserving algorithms for ordinary differential equations.

Joe, S. and F. Y. Kou (2008). Constructing Sobol sequences with better two-dimensional projections. *SIAM J. Sci. Comput 30*, 2635–2654.

Karatzas, I. and S. Shreve (1988). *Brownian Motion and Stochastic Calculus*. Springer-Verlag.

Lord, R., R. Koekkoek, and D. van Dijk (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance 10*(2), 177–194.

Lyons, T. and N. Victoir (2004). Cubature on Wiener space. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci. 460*(2041), 169–198. Stochastic analysis with applications to mathematical finance.

More, J. J., B. S. Garbow, and K. E. Hillstrom (1980). User guide for minpack-1. Technical report, Argonne National Laboratory.

Ninomiya, M. and S. Ninomiya (2009). A new higher-order weak approximation scheme for stochastic differential equations and the Runge-Kutta method. *Finance Stoch. 13*(3), 415–443.

Ninomiya, S. and N. Victoir (2008). Weak approximation of stochastic differential equations and application to derivative pricing. *Applied Mathematical Finance 15*, 107–121.

Strang, G. (1963). Accurate partial difference methods. I. Linear Cauchy problems. *Arch. Rational Mech. Anal. 12*, 392–402.

Talay, D. and L. Tubaro (1990). Expansion of the global error for numerical schemes solving stochastic differential equations. *Stochastic Anal. Appl. 8*(4), 483–509 (1991).

Chapter 5

Utility optimization and feedback-effects

Morten Karlsmark

Abstract

This paper considers a simple model for a market with one asset and two agents. One is optimizing expected utility of future wealth and the other is delta hedging derivatives written on the asset. Using the optimizing behavior of the first agent we can under simple assumptions back out the value of the delta hedgers option position. We develop the ideas first in a continuous time model, where we obtain approximations of the position, then we move on to use the LVI model of Andreasen and Huge (2011). Finally we present numerical examples where the model is applied to option data from different foreign exchange crosses to find the delta hedgers option position.

5.1 Introduction

Consider a market with one asset. This asset has a true value $(X_t)_{t\geq 0}$, a value reflecting the future cash flows and the risk averseness of the investors. But it also has a market price $(S_t)_{t\geq 0}$, which does not have to coincide with the true value in general. The difference in price exists because not all agents in the market act as investors but have other objectives, delta hedging derivatives for example.

Let us look at two different agents trading the asset, a value investor who trades according to the difference between the true value and the market price, and a delta hedger, hedging a derivative with value $F(S_t)$ to him.

Assume the demand of the value investor is $A(X_t - S_t)$ where $A(\cdot)$ is an increasing function. If the asset is cheap he will demand a positive amount and vice versa. The delta hedgers demand will be $-F'(S_t)$. This is more or less Example 3 in Schönbucher and Wilmott (2000).

We assume the asset is in fixed supply equal to a constant K. The market price S_t must then satisfy the equation

$$A(X_t - S_t) - F'(S_t) = K$$

an Itô expansion yields

$$A'(X_t - S_t)\mathrm{d}X_t - A'(X_t - S_t)\mathrm{d}S_t - F''(S_t)\mathrm{d}S_t + O(\mathrm{d}t) = 0$$

which implies

$$\mathrm{d}S_t = \frac{1}{1 + \frac{F''(S_t)}{A'(X_t - S_t)}}\mathrm{d}X_t + O(\mathrm{d}t) \tag{5.1.1}$$

where $A'(X_t - S_t) > 0$. If the delta hedger is long gamma ($F''(S_t) > 0$), volatility will be low and vice versa. Delta hedging therefore creates a feedback effect on the asset price. Note that if $F''(S_t)$ can pass $-A(X_t - S_t)$ the volatility will explode, so the model is not that well specified.

There exists a vast literature on models like this. We mention Frey and Stremme (1997), Platen and Schweizer (1998), Sircar and Papanicolaou (1998), Frey (1998) and Schönbucher and Wilmott (2000) that all consider feedback effects models and investigate the effect on stock prices and option smiles.

An implication of (5.1.1) can also be stock pinning. If delta hedgers are long straddles at certain strikes they will be long gamma in these points. The volatility in (5.1.1) will therefore go to zero around these strikes and the asset could be pinned when the straddles mature. Krishnan and Nelken (2001) and Avellaneda and Lipkin (2003) builds models to capture the effect and value options under it. Jeannin et al. (2008) builds a more complex pinning model, taking into account the results of Frey and Stremme (1997).

Empirical tests of the feedback effects hypothesis can also be found. Ni et al. (2005) tests for stock pinning in American stocks and finds that delta hedging of institutional investors contribute to stock pinning. Pearson et al. (2008) test for the impact of delta hedging besides stock pinning and finds that the delta hedging of institutional investors have a significant effect on stock price volatility.

See also Patel (2006) and Madigan (2008) for reports on feedback effects in the interest rate markets.

As noted the volatility in (5.1.1) can explode when $F''(S_t)$ can pass $-A(X_t - S_t)$. The mentioned literature have different fixes for this. Andreasen (2008) shows that this problem can be solved if the value investor acts as utility optimizer.

The normal way to specify a feedback model is to assume some behavior of the value investor and that the delta hedger is hedging a specified position. From this the dynamics of the market price can be derived. This is how the mentioned papers above build their models. Our procedure will be fundamental different from this.

We will use the intuition gained from (5.1.1) and assume the market price is a function of the underlying true value $S_t = f(X_t)$. Also we will assume the risk-neutral dynamics of S_t is of the local volatility form. Using these assumptions and a local volatility model calibrated to the market we approximate the dynamics of S_t under the real world measure. Then assuming the value investor is utility optimizing we can derive his demand from the real world dynamics of S_t. Since the value investor and the delta hedger are the only agents in the market, their demand must equal the supply of the asset and thereby we can obtain the delta hedgers demand. From this we can derive his position in the option market.

We therefore turn things upside down. Instead of assuming the delta hedgers have a given option position we will derive the position from market data.

In order to obtain our results we will need some results from utility optimizing theory. Therefore the paper is organized as follows: Section 5.2 considers classic utility optimization of terminal wealth for an agent investing in an asset following a general diffusion. Section 5.3 considers a continuous time feedback model and develops approximations to the delta hedgers position. Section 5.4 considers a discrete time discrete state space model the so called LVI model of Andreasen and Huge (2011). We show how to obtain the delta hedgers position in this model. Section 5.5 deals with numerical experiments. Here we use the results from section 5.4 on option data to derive the delta hedgers position in different foreign exchange crosses. Section 5.6 concludes.

We do all calculations in forward prices and assume the value investor invest in forwards contracts, this simplifies results. It is also possible to do a full generalization with rates and dividends.

5.2 Utility optimization

5.2.1 Exponential utility

Consider an asset with a T-forward price of S_t^T where

$$\mathrm{d}S_t^T = \mu(t, S_t^T)\mathrm{d}t + \sigma(t, S_t^T)\mathrm{d}W_t^P$$

and an investor trading in S^T who wishes to optimize his expected exponential utility at time T

$$E\left(-\exp(-\gamma V_T)\right), \quad \gamma > 0.$$

The self-financing value process V_t has the dynamics

$$\mathrm{d}V_t = \alpha_t \mathrm{d}S_t^T$$

where $(\alpha_t)_{0 \le t \le T}$ is the chosen investment strategy. As the investor buys at the money forward contracts all money transfers happen at time T. Since this

is the case the investor cannot invest in bonds, therefore we do not take them into account.

The optimal value function is given by

$$\begin{aligned} J(t,s,v) &= \sup_{\alpha} E_{t,s,v}\left(-\exp\left(-\gamma V_T\right)\right) \\ &= \exp\left(-\gamma v\right) \sup_{\alpha} E_{t,s,v}\left(-\exp\left(-\gamma \int_t^T \alpha_s \mathrm{d}S_s^T\right)\right) \end{aligned}$$

we therefore guess a solution on the form

$$J(t,s,v) = \exp\left(-\gamma v\right) j(t,s). \tag{5.2.1}$$

The HJB equation, see Björk (2004), for the problem is

$$\begin{aligned} &J_t(t,s,v) + \sup_{\alpha}\Bigg(\mu(t,s)J_s(t,s,v) + \frac{1}{2}\sigma(t,s)^2 J_{ss}(t,s,v) \\ &\quad + \alpha\mu(t,s)J_v(t,s,v) + \frac{1}{2}\alpha^2\sigma(t,s)^2 J_{vv}(t,s,v) + \alpha\sigma(t,s)^2 J_{sv}(t,s,v)\Bigg) = 0. \end{aligned}$$

Using the form of (5.2.1) and dividing by $\exp\left(-\gamma v\right)$ we get

$$\begin{aligned} &j_t(t,s) + \sup_{\alpha}\Bigg(\mu(t,s)j_s(t,s) + \frac{1}{2}\sigma(t,s)^2 j_{ss}(t,s) \\ &\quad - \gamma\alpha\mu(t,s)j(t,s) + \frac{1}{2}\gamma^2\alpha^2\sigma(t,s)^2 j(t,s) - \gamma\alpha\sigma(t,s)^2 j_s(t,s)\Bigg) = 0. \end{aligned} \tag{5.2.2}$$

As $j < 0$ the first order condition for optimal α is

$$\alpha(t,s) = \frac{\mu(t,s)}{\gamma\sigma(t,s)^2} + \frac{j_s(t,s)}{\gamma j(t,s)}. \tag{5.2.3}$$

and the second order condition is trivially satisfied since $\gamma > 0$. If $j(t,s)$ solves (5.2.2) then we know $\exp(-\gamma v)j(t,s)$ will solve the HJB equation, and then our guess will be the true solution.

5.2.2 CRRA utility

Now consider an asset with a forward price S_t^T which fulfills

$$\mathrm{d}S_t^T = \mu(t,S_t^T)S_t^T\,\mathrm{d}t + \sigma(t,S_t^T)S_t^T\,\mathrm{d}W_t^P$$

and an investor trading in S_t^T who wishes to optimize his expected CRRA utility at time T

$$E\left(\frac{V_T^\gamma}{\gamma}\right) = E\left(\frac{1}{\gamma}\exp\left(\gamma\log(V_T)\right)\right), \quad \gamma < 1.$$

The value process V_t has the dynamics

$$\begin{aligned} \mathrm{d}V_t &= \alpha_t \mathrm{d}S_t^T \\ &= \alpha_t \mu(t, S_t^T) S_t^T \mathrm{d}t + \alpha \sigma(t, S_t^T) S_t^T \mathrm{d}W_t \\ &= \beta_t \mu(t, S_t^T) V_t \mathrm{d}t + \beta_t \sigma(t, S_t^T) V_t \mathrm{d}W_t \end{aligned}$$

where $\beta_t = \frac{\alpha_t S_t^T}{V_t}$ is the chosen relative investment strategy.

Let $X_t = \log(V_t)$ then we have

$$\mathrm{d}X_t = \left(\beta_t \mu(t, S_t^T) - \frac{1}{2} \beta_t^2 \sigma(t, S_t^T)^2 \right) \mathrm{d}t + \beta_t \sigma(t, S_t^T) \mathrm{d}W_t.$$

We define the problem in terms of β and x so the optimal value function is given by

$$\begin{aligned} J(t,s,x) &= \sup_\beta E_{t,s,x} \left(\frac{\exp(\gamma X_T)}{\gamma} \right) \\ &= \frac{\exp(\gamma x)}{\gamma} \sup_\beta E_{t,s,x} \Bigg(\exp \Bigg(\gamma \int_t^T \beta_s \mu(s, S_s^T) \\ &\qquad - \frac{1}{2} \beta_s^2 \sigma(s, S_s^T)^2 \mathrm{d}s + \int_t^T \beta_s \sigma(s, S_s^T) \mathrm{d}W_s \Bigg) \Bigg) \end{aligned}$$

as in section 5.2 we guess a solution on the form

$$J(t,s,x) = \frac{\exp(\gamma x)}{\gamma} j(t,s). \tag{5.2.4}$$

The HJB equation for the problem is

$$\begin{aligned} 0 = J_t(t,s,x) + \sup_\alpha \Bigg(& \mu(t,s) s J_s(t,s,x) + \frac{1}{2} \sigma(t,s)^2 s^2 J_{ss}(t,s,x) \\ & + \left(\beta \mu(t,s) - \frac{1}{2} \beta^2 \sigma(t,s)^2 \right) J_x(t,s,x) \\ & + \frac{1}{2} \beta^2 \sigma(t,s)^2 J_{xx}(t,s,x) + \beta \sigma(t,s)^2 s J_{sx}(t,s,x) \Bigg) \end{aligned}$$

using the form of (5.2.4) and dividing with $\exp(\gamma x)$ we get:

$$\begin{aligned} 0 = j_t(t,s) + \sup_\beta \Bigg(& \mu(t,s) s j_s(t,s) + \frac{1}{2} \sigma(t,s)^2 s^2 j_{ss}(t,s) + \\ & \left(\beta \mu(t,s) - \frac{1}{2} \beta^2 \sigma(t,s)^2 \right) j(t,s) + \frac{1}{2} \gamma \beta^2 \sigma(t,s)^2 j(t,s) + \beta \sigma(t,s)^2 s j_s(t,s) \Bigg) \end{aligned}$$

as $j \neq 0$ the first order condition for optimal β is

$$\beta(t,s) = \frac{\mu(t,s)}{(1-\gamma)\sigma(t,s)^2} + \frac{j_s(t,s)s}{(1-\gamma)j(t,s)}$$

and the second order condition $(\gamma-1)\sigma(t,s)^2 j(t,s) < 0$ is satisfied when $\gamma < 1$. The optimal α is therefore given as

$$\alpha(t,s,v) = \left(\frac{\mu(t,s)s}{(1-\gamma)\sigma(t,s)^2 s^2} + \frac{j_s(t,s)}{(1-\gamma)j(t,s)} \right) v. \tag{5.2.5}$$

We can think of the v not as the wealth of the investors portfolio today but the wealth at time T. In order to start of the strategy the investor needs $v > 0$. We simply assume the investor is guaranteed some money at time T.

5.3 An approximation in a continuous time model

Consider a financial market with one traded asset, the true T-forward price follows the SDE

$$\mathrm{d}X_t^T = g(t, X_t^T)\mathrm{d}t + h(t, X_t^T)\mathrm{d}W_t^P$$

under the real world measure P. We look at a utility optimizing agent in this market. He wish to optimize the expected utility of wealth at time T, $E(u(V_T))$. Given $g(\cdot,\cdot)$ and $h(\cdot,\cdot)$ the optimization will yield a demand function $D(t, X_t^T)$.

Now assume the supply of the asset is k, we therefore obtain an equilibrium condition

$$D(t, X_t^T) = k$$

Given a utility function $u(\cdot)$ and a volatility $h(\cdot,\cdot)$, the drift $g(\cdot,\cdot)$ must satisfy a condition to make the equilibrium hold.

Now we introduce a delta hedger in the market. He has sold and bought options on the asset and delta hedges his position using the forward contract (Note that we implicitly assume the counterparties of the option contracts never trade the underlying asset). This destroys the normal equilibrium, we no longer see the true value of the asset but the market price S_t^T. The value investor will then demand $\tilde{D}(t, S_t^T)$ and the delta hedgers will demand $-F_s(t, S_t^T)$ where F is the value of his option position, this implies

$$\tilde{D}(t, S_t^T) - F_s(t, S_t^T) = k.$$

We will assume $S_t^T = f(t, X_t^T)$ where f is a smooth function with $f_x(t,x) > 0$. This seems reasonable, the delta hedger moves the price from X_t^T to S_t^T when he enters the market.

If we can find the position of the utility optimizing investor we can back out the delta hedgers delta. Thereby we can obtain $F(t,s)$ the value of the delta hedgers option position and the gamma $F_{ss}(t,s)$.

5.3.1 Specifying the dynamics of X_t^T

Assume the utility optimizing agent is the only person present in the market, and that he maximizes exponential or CRRA utility. As a start let us assume the supply of the asset is zero $k = 0$ (this could for example be a foreign exchange cross). This implies $V_t = V_0$ for all $t \in [0, T]$. He therefore knows V_T at time $t \in [0, T]$ and his expected utility will not change as a function of X_t^T. We then have $J_x(t, x, v) = 0$, using our results for the optimal strategy, see section 5.2.1 and 5.2.2, we get $g(t, x) = 0$, so X_t has zero drift.

Then assume $k \neq 0$. This is a much more complicated problem. Consider for example the exponential utility case: The PDE for $j(t, x)$ is

$$j_t(t,x) + g(t,x)j_x(t,x) + \frac{1}{2}h(t,x)^2 j_{xx}(t,x) - \gamma k g(t,x) j(t,x) + \frac{1}{2}\gamma^2 k^2 h(t,x)^2 j(t,x) - \gamma k h(t,x)^2 j_x(t,x) = 0 \quad (5.3.1)$$

and we also know

$$k = \frac{g(t,x)}{\gamma h(t,x)^2} + \frac{j_x(t,x)}{\gamma j(t,x)} \quad \Leftrightarrow \quad g(t,x) = \left(k\gamma - \frac{j_x(t,x)}{j(t,x)}\right) h(t,x)^2. \quad (5.3.2)$$

Inserting this into (5.3.1) we obtain

$$0 = j_t(t,x) - \frac{j_x(t,x)^2}{j(t,x)} h(t,x)^2 + \frac{1}{2}h(t,x)^2 j_{xx}(t,x) + j_x(t,x)\gamma k h(t,x) - \frac{1}{2}\gamma^2 k^2 h(t,x)^2 j(t,x).$$

So in order to obtain $g(t, x)$ we need to solve this nonlinear PDE and use (5.3.2). For the CRRA utility case we get an even more complicated expression. This therefore seems to be a dead end.

In the exponential utility case we could of course approximate (5.2.3) by

$$\alpha(t,x) \approx \frac{g(t,x)}{\gamma h(t,x)^2}$$

and then use

$$g(t,x) \approx k\gamma h(t,x)^2$$

as the drift. But going forwards we will assume $k = 0$, and in the empirical part only work with assets where this can be justified.

5.3.2 Specifying the dynamics for S_t^T

We have assumed $k = 0$, the dynamics of X_t^T is therefore

$$\mathrm{d}X_t^T = h(t, X_t^T)\mathrm{d}W_t^P.$$

Also we assumed $S_t^T = f(t, X_t^T)$ where $f_x(t,x) > 0$. Applying Ito's lemma we obtain

$$\begin{aligned} dS_t^T &= \left(f_t(t, X_t^T) + \frac{1}{2} f_{xx}(t, X_t^T) h(X_t^T)^2 \right) \mathrm{d}t + f_x(t, X_t^T) h(X_t^T) \mathrm{d}W_t^P \\ &= \mu(t, S_t^T)\mathrm{d}t + \sigma(t, S_t^T)\mathrm{d}W_t^P. \end{aligned} \tag{5.3.3}$$

Assume the existence of a risk neutral probability measure Q, as we are looking at the forward price we have

$$\mathrm{d}S_t^T = \sigma(t, S_t^T)\mathrm{d}W_t^Q$$

where $\sigma(t,s)$ is the local volatility function that can be obtained from European call/put prices by the Dupire equation.

By (5.3.3) we have $f_x(t,x)h(t,x) = \sigma(t, f(t,x))$. To obtain $f_{xx}(t,x)h(t,x)^2$ we differentiate $\sigma(t,s) = \sigma(t, f(t,x))$ in x

$$\frac{\partial}{\partial x}\sigma(t, f(t,x)) = \frac{\partial}{\partial x}\left(f_x(t,x)h(t,x) \right)$$

we use the chain rule on the left side

$$\sigma_s(t, f(t,x)) f_x(t,x) = f_{xx}(t,x)h(t,x) + f_x(t,x)h_x(t,x)$$

and multiply with $h(t,x)$

$$\sigma_s(t, f(t,x))\sigma(t, f(t,x)) = f_{xx}(t,x)h(t,x)^2 + \sigma(t, f(t,x))h_x(t,x)$$

which is

$$f_{xx}(t,x)h(t,x)^2 = \sigma(t, f(t,x)) \left(\sigma_s(t, f(t,x)) - h_x(t,x) \right). \tag{5.3.4}$$

If we know the local volatility function $\sigma(t,s)$ we can differentiate it and thereby we will know the first two expressions on the right hand side of (5.3.4). In order to obtain the last expression we have to specify $h(\cdot,\cdot)$, the volatility of the true price X_t^T. One could for example assume that X_t was normally ($h(t,x) = C$) or log-normally ($h(t,x) = Cx$) distributed. If we believe the feedback-effects levels out over long time horizons one could use a historic volatility as a proxy for C. Another idea would be to use an at the money volatility from the market, and then measure the feedback effects relative to this model specification.

The last thing we need to determine is $f_t(t,x)$. $f(t,x)$ can be obtained up to a function in t since (5.3.3) implies the ODE

$$f_x(t,x) = \frac{\sigma(t, f(t,x))}{h(x)}$$

for each t. Therefore we cannot determine $f_t(t,x)$ uniquely.

To get around this we could assume $f(t, x_0) = s_0$ for all t and obtain the function $f(t,x)$. But in the following we will ignore the $f_t(t,x)$ part of the drift.

5.3.3 The Delta hedger's option book

Having specified the dynamics of S_t^T under the P-measure we can obtain the optimal strategy of the value investor $\alpha(t,s)$, then we can solve for the delta hedgers option position:

Assume the delta hedger is hedging a position with the value

$$F(t,S_t^T) = E_t^Q\Bigg(\int_t^T \exp\left(-\int_t^u r(y)\mathrm{d}y\right) m\left(u,S_u^T\right)\mathrm{d}u \\ + \exp\left(-\int_t^T r(y)\mathrm{d}y\right) n\left(S_T^T\right)\Bigg)$$

where $r(y)$ is a deterministic function. The value of the claim will solve the PDE

$$F_t(t,s) + \frac{1}{2}\sigma(s)^2 F_{ss}(t,s) - r(t)F(t,s) = -m(t,s), \qquad (5.3.5)$$
$$F(T,s) = n(s)$$

In order to delta hedge this position he has to hold $-F_s(t,s)$ of the asset. The only other agent present in the market is the value investor and the market has to clear so $\alpha(t,s) = F_s(t,s)$ has to hold (we assume $k=0$). Integrating $\alpha(t,s)$ we obtain the value of the delta hedger's option position. Differentiating $\alpha(t,s)$ we obtain the gamma of the position. We can also back out the dividends $m(t,s)$ payed over each infinitesimal period.

This information could be valuable if we would like to trade options on the asset. Using the method we will know the overall positions of the delta hedgers in the market, and the gamma of their position tells us the danger zones of the volatility surface where delta hedging could become difficult.

5.3.4 Approximate optimal investment

We can use the first part of (5.2.3) or (5.2.5) as an approximation to the optimal investment strategy

$$\alpha_t \approx \frac{\mu(t,s)}{\gamma\sigma(t,s)^2} \quad \text{or} \quad \alpha_t \approx \frac{\mu(t,s)}{(1-\gamma)\sigma(t,s)^2}v. \qquad (5.3.6)$$

The expression to the right with $\gamma = 0$ is actually the optimal investment strategy for a log utility investor, see Karatzas and Shreve (1998) example 3.6.6.

$\mu(t,s)$ is approximated by

$$\mu(t,s) \approx \frac{1}{2}\sigma(t,s)\left(\sigma_s(t,s) - h_x(t,x)\right)$$

where we have to choose $h(t,x)$.

5.3.5 Approximative option position and gamma

Using (5.3.6) we can obtain an approximation of the option hedgers position. Let us use the exponential utility strategy the corresponding CRRA strategy can be found by multiplying with the wealth.

$$\begin{aligned} F(t,s) &= \int^s \frac{1}{2}\frac{\sigma_s(t,s) - h_x(t,x)}{\gamma\sigma(t,s)}\mathrm{d}s \\ &= \frac{1}{2\gamma}\left(\log(\sigma(t,s)) - \int^s \frac{h_x(t,x)}{\sigma(t,s)}\mathrm{d}s\right) \\ &= \frac{1}{2\gamma}\left(\log(\sigma(t,s)) - \int^s \frac{h_x(t,x)}{h(x)f_x(t,x)}\mathrm{d}s\right) \end{aligned}$$

we know $\frac{\mathrm{d}s}{\mathrm{d}x} = f_x(t,x)$ therefore

$$\begin{aligned} &= \frac{1}{2\gamma}\left(\log(\sigma(t,s)) - \int^x \frac{h_x(t,x)}{h(x)f_x(t,x)}f_x(t,x)\mathrm{d}x\right) \\ &= \frac{1}{2\gamma}\left(\log(\sigma(t,s)) - \log(h(t,x)) + c\right) \\ &= \frac{1}{2\gamma}\left(\log\left(\frac{\sigma(t,s)}{h(t,x)}\right) + c\right) \end{aligned} \tag{5.3.7}$$

If the local volatility of the market is higher than the true local volatility the delta hedgers have a relatively large position at this point.

To obtain the gamma we calculate

$$\begin{aligned} &\frac{\partial}{\partial s}\left(\frac{\sigma_s(t,s) - h_x(t,x)}{2\gamma\sigma(t,s)}\right) \\ &\quad= \frac{1}{2\gamma}\left(\frac{\sigma_{ss}(t,s) - h_{xs}(t,x)}{\sigma(t,s)} - \frac{\sigma_s(t,s)\left(\sigma_s(t,s) - h_x(t,x)\right)}{\sigma(t,s)^2}\right) \\ &\quad= \frac{1}{2\gamma}\left(\frac{\sigma_{ss}(t,s) - h_{xx}(t,x)f_x(t,x)^{-1}}{\sigma(t,s)} - \frac{\sigma_s(t,s)\left(\sigma_s(t,s) - h_x(t,x)\right)}{\sigma(t,s)^2}\right) \\ &\quad= \frac{1}{2\gamma}\Bigg(\frac{\sigma_{ss}(t,s)\sigma(t,s) - h_{xx}(t,x)f_x(t,x)^{-1}\sigma(t,s)}{\sigma(t,s)^2} - \\ &\qquad\qquad \frac{\sigma_s(t,s)\left(\sigma_s(t,s) - h_x(t,x)\right)}{\sigma(t,s)^2}\Bigg) \\ &\quad= \frac{1}{2\gamma}\left(\frac{\sigma_{ss}(t,s)\sigma(t,s) - h_{xx}(t,x)h(t,x) - \sigma_s(t,s)\left(\sigma_s(t,s) - h_x(t,x)\right)}{\sigma(t,s)^2}\right). \end{aligned}$$

We see that the gamma depends on the differences in the first and second order derivatives of the local volatility of the market and true local volatility function.

5.4 A discrete model: LVI

Instead of working with a continuous time model we can work with a discrete model. The so called LVI model, see Andreasen and Huge (2011), can be calibrated to a market of European call and put options and we can also easily solve the investors optimization problem in the model.

5.4.1 The LVI model

The price of a derivative in our continuous time model (neglecting discounting) is the solution to the PDE

$$0 = v_t(t,s) + \frac{1}{2}\sigma(t,s)^2 v_{ss}(t,s)$$

We approximate the PDE by an implicit finite differences scheme

$$0 = \frac{v_{t_{i+1}} - v_{t_i}}{\Delta t_i} + A_i v_{t_i} \tag{5.4.1}$$

here A_i is a finite difference matrix containing both the function $\sigma(t,s)^2$ and the second order finite difference:

$$A_i = \begin{pmatrix} 0 & 0 & \cdots & & \\ \frac{\sigma(t_i,s_1)^2}{2\Delta s^2} & -\frac{\sigma(t_i,s_1)^2}{\Delta s^2} & \frac{\sigma(t_i,s_1)^2}{2\Delta s^2} & 0 & \cdots \\ \ddots & \ddots & \ddots & \ddots & \ddots \\ \cdots & 0 & \frac{\sigma(t_i,s_{n-1})^2}{2\Delta s^2} & -\frac{\sigma(t_i,s_{n-1})^2}{\Delta s^2} & \frac{\sigma(t_i,s_{n-1})^2}{2\Delta s^2} \\ & & \cdots & 0 & 0 \end{pmatrix}$$

Note that we have set the second order derivative equal to zero on the boundary.

Rearranging (5.4.1), we obtain

$$(I - \Delta t_i A_i) v_{t_i} = v_{t_{i+1}} \Leftrightarrow v_{t_i} = (I - \Delta t_i A_i)^{-1} v_{t_{i+1}}$$

where I is the identity matrix. Define $C = (I - \Delta t_i A_i)^{-1}$. As $(I - \Delta t_i A_i)$ is an M-Matrix, see Horn and Johnson (1991), we know that $C_{ij} \geq 0$ and since $(I - \Delta t_i A_i)\mathbf{1} = \mathbf{1}$ we also have $C\mathbf{1} = \mathbf{1}$ where $\mathbf{1} = (1, \ldots, 1)'$. Therefore C can be interpreted as the transition probabilities of the discrete system (C_{ij} is is the transition probability from state s_i to state s_j). We also get $(I - \Delta t_i A_i)s = s$ so the process defined by C will be a martingale. By specifying the finite difference scheme we have in fact specified an arbitrage free discrete state space discrete time model for the forward under the Equivalent Martingale Measure Q.

The model can be calibrated to any European option surface yielding a discrete local volatility function $\sigma(t,s)$.

5.4.2 Specification of the P-measure in the discrete model

As in the continuous model we will assume the supply of the asset is zero, therefore we would like the true price X_t^T to be a martingale under the P-measure. We assume the transition probabilities for X_t^T from time t_i to t_{i+1} under P can be written as

$$\left(I - \frac{\Delta t_i}{2} h(t_i, x)^2 \delta_{xx}\right)^{-1}$$

Here δ_{xx} is the discrete second order finite difference operator.

Since x is a linear function of x we have

$$\left(I - \frac{\Delta t_i}{2} h(t_i, x)^2 \delta_{xx}\right) x = x$$

so X_t^T will be a martingale under the P measure.

In the same manner we assume that the transition probabilities for S_t^T from t_i to t_{i+1} under Q can be written on the form

$$\left(I - \frac{\Delta t_i}{2} \sigma(t_i, s)^2 \delta_{ss}\right)^{-1}.$$

$\sigma(t_i, s)$ is found by calibrating the model to the market, see section 5.4.1.

Under the P-measure we specify the transition probabilities for S_t^T as

$$\left(I - \frac{\Delta t_i}{2} \sigma(t_i, s)^2 \delta_{ss} - \Delta t_i \mu(t_i, s) \delta_s\right)^{-1}.$$

Here we might have to use winding in order to ensure positive transition probabilities, see for example Karlsmark (2013).

We have

$$\left(I - \frac{\Delta t_i}{2} \sigma(t_i, s)^2 \delta_{ss} - \Delta t_i \mu(t_i, s) \delta_s\right) s = s - \Delta t_i \mu(t_i, s)$$

since $\delta_s s = 1$, so S_t^T has a drift under the P-measure.

Then we assume $s = f(x)$ where f is an increasing function. f is chosen like this because our grid in s and x is fixed for all times. The transition probability from s_i to s_j is the same as the transition probability from x_i to x_j so we obtain the system

$$\left(I - \frac{\Delta t_i}{2} \sigma(t_i, s)^2 \delta_{ss} - \Delta t_i \mu(t_i, s) \delta_s\right)^{-1} = \left(I - \frac{\Delta t_i}{2} h(t_i, x)^2 \delta_{xx}\right)^{-1} \Leftrightarrow$$
$$\frac{1}{2} \sigma(t_i, s)^2 \delta_{ss} + \mu(t_i, s) \delta_s = \frac{1}{2} h(t_i, x)^2 \delta_{xx}$$

Unfortunately this system cannot be solved, we need a different $\mu(t_i, s)$ not only for each row, but also for each column. As an approximation we solve the

equation system when the finite difference operators are replaced by differential operators:

$$\frac{1}{2}\sigma(t_i,s)^2\frac{\partial^2}{\partial s^2}+\mu(t_i,s)\frac{\partial}{\partial s}=\frac{1}{2}h(t_i,x)^2\frac{\partial^2}{\partial x^2}\Leftrightarrow$$
$$\frac{1}{2}\sigma(t_i,s)^2\frac{\partial^2}{\partial s^2}+\mu(t_i,s)\frac{\partial}{\partial s}=\frac{1}{2}h(t_i,x)^2\left(f'(x)^2\frac{\partial^2}{\partial s^2}+f''(x)\frac{\partial}{\partial s}\right)$$

This equation has to hold for any function we multiply onto both sides. If we multiply a linear function of s we get

$$\mu(t_i,s)=\frac{1}{2}h(t_i,x)^2f''(x) \tag{5.4.2}$$

therefore we also have

$$\frac{1}{2}\sigma(t_i,s)^2=\frac{1}{2}h(t_i,x)^2f'(x)^2. \tag{5.4.3}$$

Differentiating (5.4.3) in x we obtain

$$\sigma(t_i,s)\sigma_s(t_i,s)f'(x)=h(t_i,x)h_x(t_i,x)f'(x)^2+h(t_i,x)^2f'(x)f''(x)$$

which is

$$f''(x)h(t_i,x)^2=\sigma(t_i,s)\sigma_s(t_i,s)-h(t_i,x)h_x(t_i,x)f'(x)$$

and we get

$$\mu(t_i,s)=\frac{1}{2}\sigma(t_i,s)\left(\sigma_s(t_i,s)-h_x(t_i,x)\right).$$

This looks like the result in the continuous case. But we have chosen f to be a function only of x therefore we have no f_t part in this expression.

The LVI model provides us with a discrete function $\sigma(t_i,s)$, we then approximate $\sigma_s(t_i,s)$ with finite differences. Then we need to specify the $h_x(t_i,x)$ function, this will be dealt with in the specific examples of section 5.5. On the boundary of our finite difference matrix we will set both the first and second order derivative equal to zero.

We have then specified the LVI version of the model, and we can solve for the optimal investment strategy of the value investor.

5.4.3 Utility optimization

We follow Andreasen (2008). Consider an investor who holds the initial wealth V_0 and invest in $S_t^{t_n}$. He wants to maximize his expected utility

$$\sup_{\{\alpha_j\}_{0\leq j\leq n-1}} E\left(u\left(V_0+\sum_{j=0}^{n-1}\alpha_j\left(S_{t_j}^{t_n},V_{t_j}\right)\left(S_{t_{j+1}}^{t_n}-S_{t_j}^{t_n}\right)\right)\right)$$

In full generality this problem is two-dimensional since α depends on S^{t_n} and V. We can get around this if we work with the exponential or the CRRA utility function.

Let us consider the exponential utility function $u(x) = -\exp(-\gamma x)$ then

$$\begin{aligned}
&\sup_{\{\alpha_j\}_{0\le j\le n-1}} \mathbb{E}\left(-\exp\left(-\gamma V_{t_n}\right)\right) \\
&= \sup_{\{\alpha_j\}_{0\le j\le n-1}} \mathbb{E}\left(-\exp(-\gamma V_{t_{n-1}})\exp\left(-\gamma\alpha_{n-1}\left(S^{t_n}_{t_{n-1}}, V_{t_{n-1}}\right)\left(S^{t_n}_{t_n} - S^{t_n}_{t_{n-1}}\right)\right)\right) \\
&= \sup_{\{\alpha_j\}_{0\le j\le n-2}} \mathbb{E}\Bigg(\exp(-\gamma V_{t_{n-1}}) \\
&\qquad\qquad \sup_{\alpha_{n-1}} \mathbb{E}_{t_{n-1}}\left(-\exp\left(-\gamma\alpha_{n-1}\left(S^{t_n}_{t_{n-1}}, V_{t_{n-1}}\right)\left(S^{t_n}_{t_n} - S^{t_n}_{t_{n-1}}\right)\right)\right)\Bigg).
\end{aligned}$$

Now consider

$$\sup_{\alpha_{n-1}} \mathbb{E}_{t_{n-1}}\left(-\exp\left(-\gamma\alpha_{n-1}\left(S^{t_n}_{t_{n-1}}, V_{t_{n-1}}\right)\left(S^{t_n}_{t_n} - S^{t_n}_{t_{n-1}}\right)\right)\right)$$

and note that the optimal α_{n-1} does not depend on $V_{t_{n-1}}$ since $(S^{t_n}_{t_i})_{0\le i\le n}$ is Markov. We therefore define

$$K_{n-1}\left(S^{t_n}_{t_{n-1}}\right) = \sup_{\alpha_{n-1}} \mathbb{E}_{t_{n-1}}\left(-\exp\left(-\gamma\alpha_{n-1}\left(S^{t_n}_{t_{n-1}}\right)\left(S^{t_n}_{t_n} - S^{t_n}_{t_{n-1}}\right)\right)\right).$$

Continuing backwards we also define

$$K_i\left(S^{t_n}_{t_i}\right) = \sup_{\alpha_i} \mathbb{E}_{t_i}\left(\exp\left(-\gamma\alpha_i\left(S^{t_n}_{t_i}\right)\left(S^{t_n}_{t_{i+1}} - S^{t_n}_{t_i}\right)\right)K_{i+1}\left(S^{t_n}_{t_{i+1}}\right)\right)$$

as the optimal α_i does not depend on V_{t_i}. By definition we have

$$K_i\left(S^{t_n}_{t_i}\right) = \sup_{\{\alpha_j\}_{i\le j\le n-1}} \mathbb{E}_{t_i}\left(-\exp\left(-\gamma\sum_{j=i}^{n-1}\alpha_j\left(S^{t_n}_{t_j}, V_{t_j}\right)\left(S^{t_n}_{t_{j+1}} - S^{t_n}_{t_j}\right)\right)\right).$$

Assume $K_{i+1}\left(S^{t_n}_{t_{i+1}}\right)$ has been obtained. To find α_i we differentiate $K_i(S_{t_i})$ in α_i and obtain the equation

$$\mathbb{E}_{t_i}\left(\frac{\partial}{\partial\alpha_i}\exp\left(-\gamma\alpha_i\left(S^{t_n}_{t_{i+1}} - S^{t_n}_{t_i}\right)\right)K_{i+1}\left(S^{t_n}_{t_{i+1}}\right)\right) = 0. \tag{5.4.4}$$

We can write this as

$$\mathbb{E}_{t_i}\left(-\gamma\left(S^{t_n}_{t_{i+1}} - S^{t_n}_{t_i}\right)\exp\left(-\gamma\alpha_i\left(S^{t_n}_{t_{i+1}} - S^{t_n}_{t_i}\right)\right)K_{i+1}\left(S^{t_n}_{t_{i+1}}\right)\right) = 0$$

which gives us

$$\mathbb{E}_{t_i}\left(S_{t_{i+1}}^{t_n}\exp\left(-\gamma\alpha_i S_{t_{i+1}}^{t_n}\right)K_{i+1}\left(S_{t_{i+1}}^{t_n}\right)\right) - S_{t_i}^{t_n}\mathbb{E}_{t_i}\left(\exp\left(-\gamma\alpha_i S_{t_{i+1}}^{t_n}\right)K_{i+1}\left(S_{t_{i+1}}^{t_n}\right)\right)=0.$$

Define

$$L(t_i,\alpha)=\mathbb{E}_{t_i}\left(S_{t_{i+1}}^{t_n}\exp\left(-\gamma\alpha S_{t_{i+1}}^{t_n}\right)K_{i+1}\left(S_{t_{i+1}}^{t_n}\right)\right)$$

and

$$M(t_i,\alpha)=\mathbb{E}_{t_i}\left(\exp\left(-\gamma\alpha S_{t_{i+1}}^{t_n}\right)K_{i+1}\left(S_{t_{i+1}}^{t_n}\right)\right).$$

For a fixed α, $L(t_i,\alpha)$ and $M(t_i,\alpha)$ are the solutions to

$$\begin{aligned}(I-\Delta t_i A_i)\,L(t_i,\alpha)&=s\exp\left(-\gamma\alpha s\right)K_{i+1}\left(s\right)\\(I-\Delta t_i A_i)\,M(t_i,\alpha)&=\exp\left(-\gamma\alpha s\right)K_{i+1}\left(s\right)\end{aligned}$$

where s is the grid of stock prices and $A_i=\frac{1}{2}\sigma(t_i,s)^2\delta_{ss}+\mu(t_i,s)\delta_s$ is our finite difference matrix for the period $(t_i,t_{i+1}]$ under the P-measure.

We choose N constants $(\alpha_j)_{1\leq j\leq N}$ and solve for L and M for each α_j. We obtain

$$L(t_i,\alpha_j)_k-s_k M(t_i,\alpha_j)_k\quad j=1,\ldots,N \tag{5.4.5}$$

for each grid point s_k, and we wish to find α that makes (5.4.5) equal zero. By a search in $(\alpha_j)_{1\leq j\leq N}$ we obtain α_l and α_{l+1} such that $\alpha_l\leq\alpha\leq\alpha_{l+1}$. The optimal α is then approximated by linear interpolation between α_l and α_{l+1} given the slope of the function (5.4.5).

The procedure described here is more or less "piecewise constant policy timestepping" with an extra interpolation of α, see Forsyth and Labahn (2008) and the references therein.

In the CRRA utility case we can do the following for the last time step:

$$\begin{aligned}&\sup_{\alpha_{n-1}}\mathbb{E}_{n-1}\left(\frac{1}{\gamma}\left(V_{t_{n-1}}+\alpha_{n-1}\left(S_{t_{n-1}}^{t_n},V_{t_{n-1}}\right)\left(S_{t_n}^{t_n}-S_{t_{n-1}}^{t_n}\right)\right)^{\gamma}\right)\\&=\frac{1}{\gamma}V_{t_{n-1}}^{\gamma}\sup_{\alpha_{n-1}}\mathbb{E}_{n-1}\left(\left(1+\frac{\alpha_{n-1}\left(S_{t_{n-1}}^{t_n},V_{t_{n-1}}\right)}{V_{t_{n-1}}}\left(S_{t_n}^{t_n}-S_{t_{n-1}}^{t_n}\right)\right)^{\gamma}\right)\\&=\frac{1}{\gamma}V_{t_{n-1}}^{\gamma}\sup_{\beta_{n-1}}\mathbb{E}_{n-1}\left(\left(1+\beta_{n-1}\left(S_{t_{n-1}}^{t_n},V_{t_{n-1}}\right)\left(S_{t_n}^{t_n}-St_{n-1}^{t_n}\right)\right)^{\gamma}\right).\end{aligned}$$

Note that the optimal β_{n-1} will not depend on $V_{t_{n-1}}$ and we can solve for it using the same procedure as above. Doing this recursively we obtain $\beta_i(S_{t_i})$ for $i=0,\ldots,n-1$ and we have $\alpha_i(S_{t_i},V_{t_i})=\beta_i(S_{t_i})V_{t_i}$. If we know the initial wealth of the value investor we can calculate how much he invests in the asset at time t_0.

5.4.4 The delta hedger

In the continuous time model we know the delta hedger will hold $-F_s(t,s)$ of the asset where $F(t,s)$ is the value of his option portfolio. As an approximation we will assume the same thing is true in our discrete model. Also as in section 5.3.3 we assume the delta hedger's position in the asset is the opposite of the value investors position.

Therefore we can obtain the value of the delta hedgers position by integrating $\alpha(t,s)$ in s for each t, this is done numerically by the trapezoidal rule. We can also obtain the gamma by differentiating $\alpha(t,s)$ in s, this can be done by finite differences.

When we have the value of the delta hedgers position we can in fact also compute the future dividend payments by him. This is done using the discrete version of (5.3.5) to get

$$M(t_i, t_{i+1}) = \frac{1}{\Delta t_i}\left((I - \Delta t_i A_i)F(t_i) - F(t_{i+1})\right).$$

A_i will now also include discounting.

5.5 Numerical examples

5.5.1 The Data

We will use our discrete model in the foreign exchange market. The foreign exchange market is chosen because investors are not naturally long or short a foreign exchange cross. We can therefore assume the supply of the underlying is zero when no delta hedgers are present. This assumption will for example clearly be wrong in the equity markets.

We consider four FX crosses: USD/JPY, EUR/USD, USD/GBP and EUR/CHF for the date March 19, 2013. Data where gathered from Bloomberg using the Excel function "BDH". We use the option quotes ATM, RR5, RR10, RR15, RR25, RR35, BF5, BF10, BF15, BF25, and BF35 for the maturities 1W, 2W, 3W, 1M, 2M, 3M, 6M, 1Y, 18M, 2Y, 3Y, and 5Y.

The FX forward curves were built using FX forward quotes from Bloomberg, we use the quotes SPOT, 1W, 2W, 3W, 1M, 2M, 3M, 4M, 5M, 6M, 9M, 12M, 18M, 2Y, 3Y, 4Y, and 5Y.

5.5.2 Choosing the underlying model

A natural way to specify the underlying model in the FX markets is to choose a log-normal model ie $h(t,x) = Kx$ in the continuous time case. But the LVI model is not a continuous time model, the choice $h(t,x) = Kx$ would not correspond to a log-normal model.

All our option quotes are given as Black volatilities, in order to obtain an h function that corresponds to a log-normal model we will do the following:

For each maturity we set the Black volatilities of all options equal to the ATM Black volatility. Thereby we obtain a flat Black smile for each maturity. Then we calibrate the LVI model to these quotes and obtain a function $h(t,x)$, which can be differentiated by finite difference to obtain $h_x(t,x)$.

Let us then take a look at the approximative strategy from section 5.3

$$\alpha_t \approx \frac{\sigma_s(t,s) - h_x(t,x)}{\gamma\sigma(t,s)}.$$

We can write $\sigma(t,s)$ and $h(t,x)$ on the Black form instead and we see that

$$\sigma_s(t,s) - h_x(t,x) = \sigma_s^{Black}(t,s)s + \sigma^{Black}(t,s) - \left(h_x^{Black}(t,x)x + h^{Black}(t,x)\right)$$

Our LVI parametrization does a flat extrapolation of the Black volatility. Therefore the difference $\sigma^{Black}(t,s) - h^{Black}(t,x)$ will make the delta hedgers position increase or decrease linearly for very high and low strikes, as we integrate α_t in the s direction.

To counteract this we choose to remove $\sigma^{Black}(t,s) - h^{Black}(t,x)$ from $\mu(t,s)$ such that the trading strategy becomes

$$\alpha_t \approx \frac{\sigma_s^{Black}(t,s)s - h_x^{Black}(t,x)x}{\gamma\sigma^{Black}(t,s)s}.$$

We then capture more naturally the important parts of the delta hedgers position.

The choice of $h(t,x)$ is of course very arbitrary. If we had another view on the true smile we could calibrate the LVI model to this and obtain another result. One could for example set $h(t,x)$ equal to yesterdays local volatility function $\sigma(t,s)$. Doing that we would more or less obtain the delta hedgers position relative to yesterdays position.

5.5.3 Grid specification

In order to interpolate the local volatility function we use a log transformed finite difference scheme:

$$\left(I - \frac{\Delta t}{2}\sigma^{Black}(t,\exp(x))^2(\delta_{xx} - \delta_x)\right) v_i(\exp(x)) = v_{i+1}(\exp(x))$$

where $s = \exp(x)$. $\exp(x)$ will not be a martingale in this scheme, but the obtained local volatility function will be nicer than the one obtained in a normal scheme.

For the optimization we will also use a log-transformed scheme. As we include a drift for the asset, the martingale property will not be important.

Note that in both of the above cases we still obtain non-negative transition densities if the matrix $(I - \Delta tA)$ (A is our finite difference matrix) is an M-matrix, see Horn and Johnson (1991). This will be true if our finite difference

grid is dense enough or if we use winding. In the examples below the finite difference grid will be dense enough so we will not use winding.

Under the risk neutral measure the spot of the FX cross will have a drift $\mu(s)$. We use the time 0 forward F_t^0 to interpolate the local volatility function. Here we define

$$F_t^0 = S_t \exp\left(-\int_0^t \mu(s)\mathrm{d}s\right).$$

So our local volatility function will actually be a function of F_t^0. To get a function of our time T forward we add $\log\left(\frac{F_0^T}{F_0^0}\right)$ to all grid points. Then we can do the portfolio optimization in the forward $(F_t^T)_{0\leq t\leq T}$.

The grid of the log-transformed scheme will span the interval $[\log(F_0^T) - 5\sigma T, \log(F_0^T)+5\sigma T]$, where σ is the Black volatility for the longest ATM option and T is the maturity for the longest options. It will contain 500 grid points and we will use 1000 different test α's when we search for the optimal α.

5.5.4 Delta hedgers position on March 19 2013

We apply the method to the four option smiles from March 19, 2013. Let us start by looking at the USD/JPY example. We plot the local volatilities obtained by the LVI model divided with the local volatilities from our true model (the log-normal model).

Then we plot the delta hedgers position using the same data, this can be seen in figure 5.2. In the figure we start at time 0, the graph below is the value of the delta hedgers position at time 0 as a function of the underlying. We end at time 3 this is the value of the delta hedgers position 3 years from now, if the delta hedgers do not take on new option positions.

Note first of all how much figure 5.1 and 5.2 looks like each other, the approximation (5.3.7) seems to hold quite well even in the LVI model.

Then note how the volatility in figure 5.1 follows the position in figure 5.2. Left of ATM the position goes from flat to downward sloping this implies that the delta hedger has a negative gamma which by (5.1.1) implies a high volatility. Around ATM the position goes from downward sloping to upward sloping therefore the gamma will be positive and the volatility will be low. Right of ATM the position goes from upward sloping to flat this implies a negative gamma and therefore a high volatility. Everything fits with (5.1.1).

Next we plot the delta hedgers position for the other crosses these can be seen in figure 5.3 to 5.5.

The level of the position is of course relative. We have used the exponential utility function to obtain the results, since we have no information about the wealth of the value investor. Also we do not know the overall level of the position (the c in (5.3.7)) therefore we have just used 0.

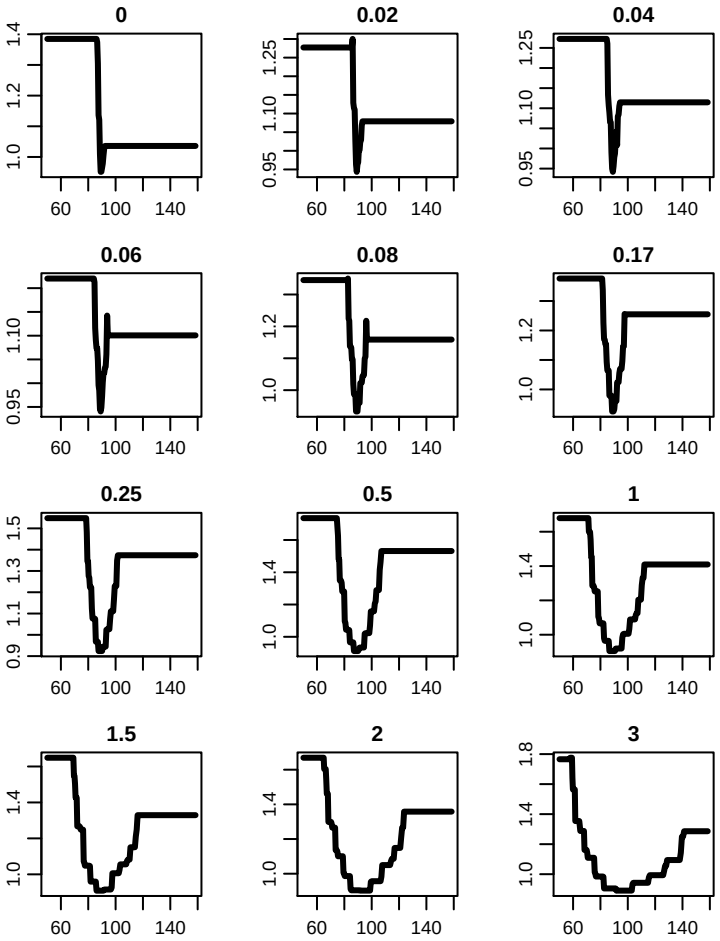

Figure 5.1: The local volatility of the market divided with the true local volatility for USD/JPY on March 19, 2013

5.6 Conclusion

This paper has considered a method to obtain the value of the delta hedger's option position in a market. The results rely on a number of assumptions, among them that the value investor and delta hedgers are the only agents present in the market. We also need to specify the dynamics of the underlying when no delta hedgers are present. The numerical results will therefore depend heavily on these assumptions, but the results can give an indication of the delta hedger's position. The method is therefore a step towards a practical use of the feedback effects models developed by Frey and Stremme (1997) and others. Future research could include models with stochastic volatility or jumps as well as applying the model to equity markets where we cannot make the assumption that the supply of the underlying is zero.

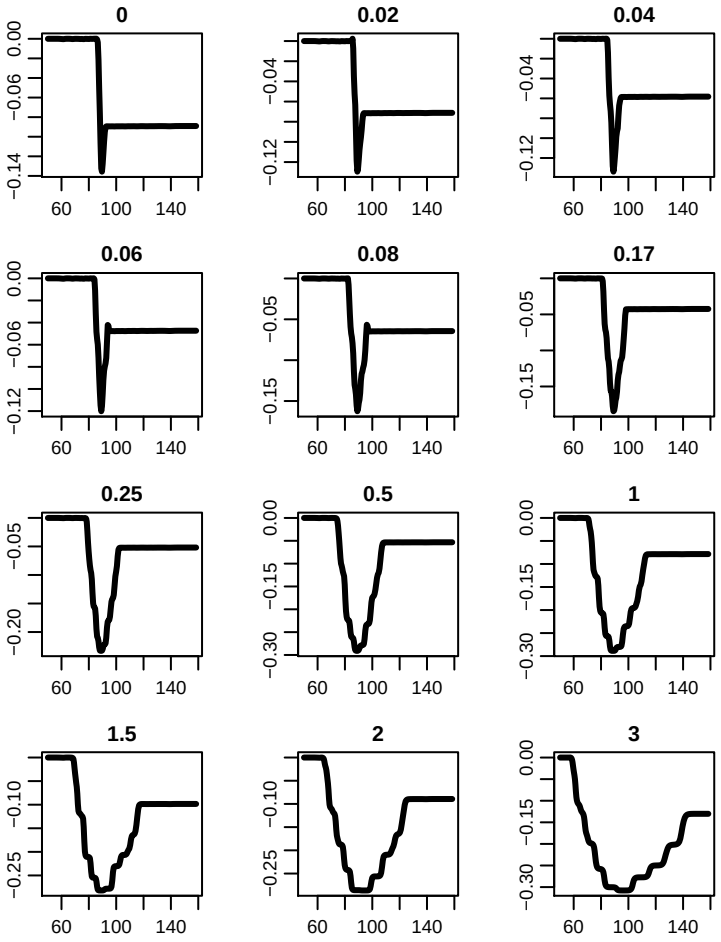

Figure 5.2: The value of the delta hedger position in USD/JPY on March 19, 2013

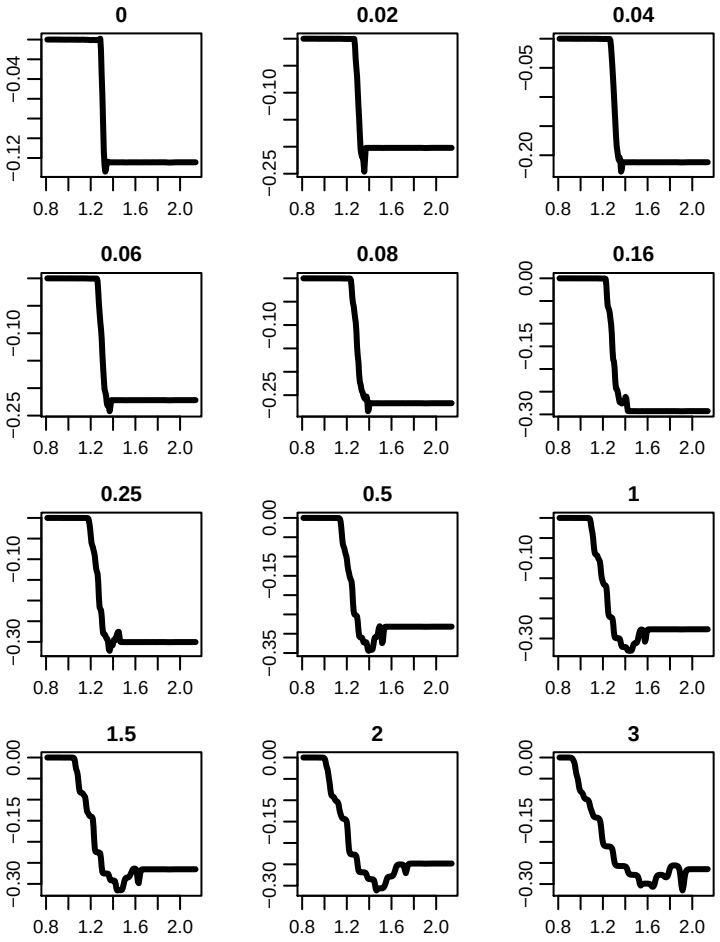

Figure 5.3: The value of the delta hedger position in EUR/USD on March 19, 2013

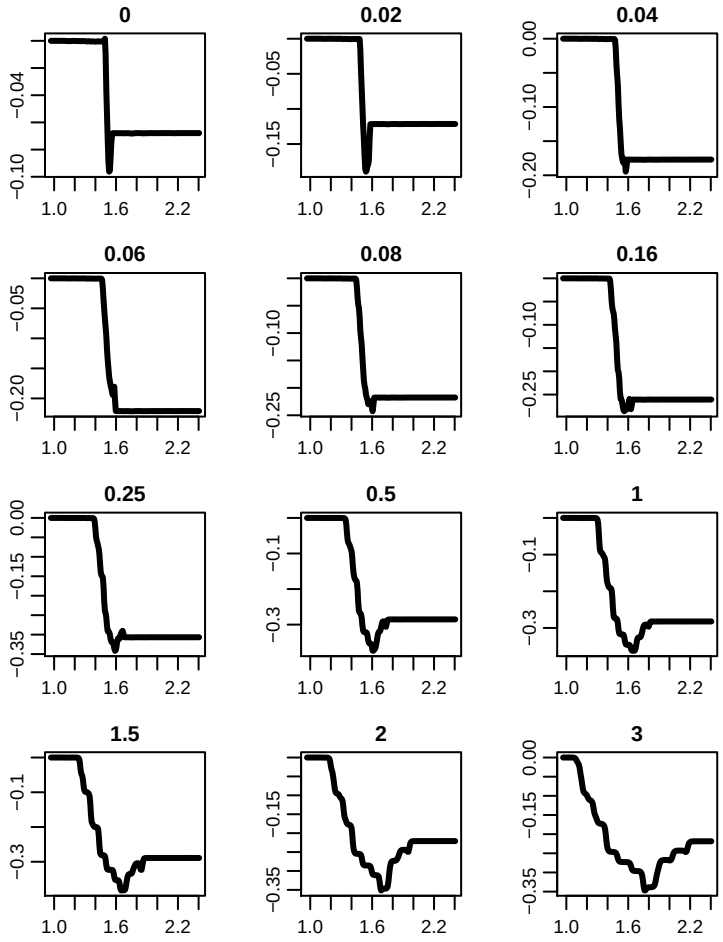

Figure 5.4: The value of the delta hedger position in GBP/USD on March 19, 2013

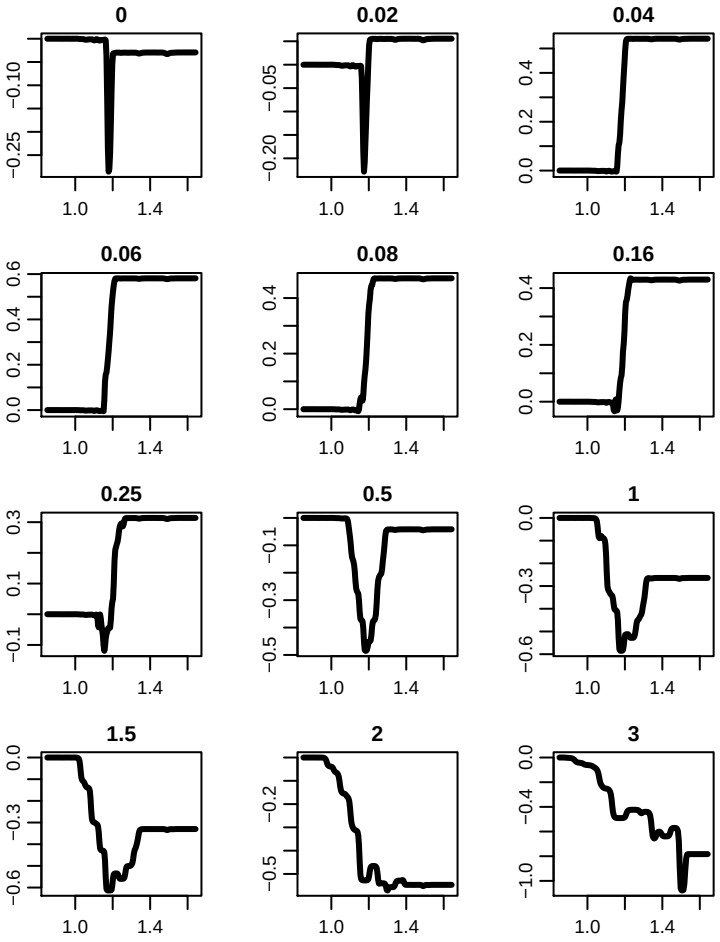

Figure 5.5: The value of the delta hedger position in EUR/CHF on March 19, 2013

Bibliography

Andreasen, J. (2008). Soft markets: Feedback effects from dynamic hedging. Presentation, Warwick University.

Andreasen, J. and B. Huge (2011, March). Volatility interpolation. *Risk Magazine 24*(3), 76–79.

Avellaneda, M. and M. D. Lipkin (2003). A market-induced mechanism for stock pinning. *Quantitative finance 3*, 417–425.

Björk, T. (2004). *Arbitrage theory in continuous time* (Second ed.). Oxford University Press.

Forsyth, P. A. and G. Labahn (2008). Numerical methods for controlled hamilton-jacobi-bellman pdes in finance. *Journal of Computational Finance 11*, 1–44.

Frey, R. (1998). Perfect option hedging for a large trader. *Finance and stochastics 2*, 115–141.

Frey, R. and A. Stremme (1997). Market volatility and feedback effects from dynamic hedging. *Mathematical Finance 7*(4), 351–374.

Horn, R. A. and C. R. Johnson (1991). *Topics in Matrix Analysis*. Cambridge University Press.

Jeannin, M., G. Iori, and D. Samuel (2008). Modeling stock pinning. *Quantitative Finance 8*(8), 823–831.

Karatzas, I. and S. E. Shreve (1998). *Methods of Mathematical Finance*. Springer.

Karlsmark, M. (2013). A numerical scheme for perpetual claims. Working paper.

Krishnan, H. and I. Nelken (2001, December). The effect of stock pinning upon option prices. *RISK*, 17–20.

Madigan, P. (2008). The rates escape. *Risk Magazine*.

Ni, S. X., N. D. Pearson, and A. M. Poteshman (2005). Stock price clustering on option expiration dates. *Journal of Financial Economics 78*, 49–87.

Patel, N. (2006). The gamma trap. *Risk Magazine*.

Pearson, N. D., A. M. Poteshman, and J. White (2008). Does option trading have a pervasive impact on underlying stock prices? Working paper.

Platen, E. and M. Schweizer (1998). On feedback effects from hedging derivatives. *Mathematical Finance 8*(1), 67–84.

Schönbucher, P. J. and P. Wilmott (2000). The feedback effect of hedging in illiquid markets. *SIAM J. Appl. Math. 61*(1), 232–272.

Sircar, K. R. and G. Papanicolaou (1998). General Black-Scholes models accounting for increased market volatility from hedging strategies. *Applied Mathematical Finance 5*, 45–82.