Jethwa_Jay_01513119

by Jay Jethwa

Submission date: 05-Sep-2022 06:27PM (UTC+0100)
Submission ID: 185719620
File name: Jethwa_Jay_01513119.pdf (943.26K)
Word count: 16159
Character count: 73850

Imperial College London

Imperial College London

Department of Mathematics

Uncertain Volatility Model with Applications to Cliquets

Author: Jay Jethwa (CID: 01513119)

A thesis submitted for the degree of

MSc in Mathematics and Finance, 2021-2022

Declaration

The work contained in this thesis is my own work unless otherwise stated.

Acknowledgements

I would like to express my gratefulness to my thesis supervisors Dr Jack Jacquier and Dr Alexandre Pannier for their guidance.

I would also like to thank my family for all of their motivation and support throughout my studies.

Abstract

Pricing and hedging derivatives is a classical problem within Mathematical Finance. A useful tool for tackling these problems are Monte Carlo methods, which generally require some underlying model. Often it is not possible to calibrate these models to the market accurately — the Uncertain Volatility Model aims to circumvent this issue.

Within this thesis we begin by introducing the Uncertain Volatility Model, and discuss its application to the super-replicating problem. We then explore various numerical methods, in particular we develop a robust pricing method for a variety of cliquets.

Keywords: Uncertain Volatility Model, Black-Scholes-Barenblatt, Super-Replicating, Cliquets

Contents

List of Figures

Introduction

Derivative contracts are financial products that usually guarantee some stream of cash flows contingent on one or more underlying assets. They are extremely popular amongst investors as they allow one to take a more specialised view on the market and hedge (or leverage) specific risks. For example, an in-the-money *butterfly* option allows an investor to speculate on the magnitude of the return of a single underlying, irrespective of its direction. A more complex product called a *cliquet*, provides an investor with protection against extreme moves in the return of an asset, usually paired with some global feature expressing a more specific view. This is an example of an *exotic* product and is not usually liquid (or existent) on an exchange.

Financial institutions that wish to sell complex products to investors often need to develop models for pricing these derivatives. The price they choose must be large enough such that their hedges cover any losses they incur from their obligations of the contract, but the price must also be small enough such that they remain competitive amongst other institutions. One of the most important components of these models is the dynamics that the volatility is assumed to follow. There are many variations; in Black-Scholes the volatility is assumed to be fixed over the time horizon, local volatility models assume the volatility is a deterministic function of time and the spot, stochastic volatility models assume that volatility follows some random process such as in the Heston model. All of these models, before being used for pricing must be calibrated to the underlying asset; this involves calibrating the volatility surface produced by the model to the volatility surface implied by the market. The market implied volatility surface is usually obtained by finding the Black-Scholes implied volatilities for a variety of **liquid** call/put options on the underlying. For more information on this matter we refer the reader to "The Volatility Surface : A Practitioner's Guide" by Jim Gatheral[7].

So it is clear that if we are to take this common approach to derivative pricing, we need data to calibrate our volatility surface. However, for many assets there may be a lack of liquidity in exchange traded derivatives. In some cases, where the underlying process we consider is a proprietary strategy, there may be no data to calibrate our volatility surface. A solution, proposed by Avellaneda et al [1] and Lyons [11], is to instead consider a family of models where the volatility is assumed to lie in some range of values. In this setting, we naturally obtain a family of prices of which we desire the maximum, so as to hedge the 'worst' possible outcome. The selection of the bounds are normally considered to be deterministic (possibly time varying), however there has been some interesting research into stochastic bounds by Fouque and Ning [6]. There have also been studies of the extreme cases; Martini [12] studied the case where the lower bound is assumed to be zero, Leblanc and Martini [10] studied the case where volatility is unbounded.

This paper is structured as follows; in Chapter 1 we introduce the core ideas and results when pricing derivatives with uncertain volatility ([1], [11]). In Chapter 2 we explore the link between the maximal arbitrage-free price produced by the family of models and the cheapest super-replicating strategy which is of practical importance ([3]). In Chapter 3 we introduce some new theoretical results on so called 'zero gamma boundaries' which become useful in our numerical methods. In Chapters 4 and 5 we discuss some numerical methods presented in literature ([1], [8]) and develop them to build a new robust pricing method for a variety of cliquets under model uncertainty.

Chapter 1

Introduction to Uncertain Volatility

1.1 Preliminaries

Before proceeding with the main content of the thesis, we cover a brief overview of some basic derivative pricing, outlining the key concepts and problems that are of particular interest. We then discuss super-replicating strategies, and why they are of practical importance. Finally, we define the Greeks and build some intuition around them. The following preliminaries are based on "PDE and Martingale Methods in Option Pricing" by Andrea Pascucci [13].

Suppose in our market we have a single asset, whose price process is denoted by $(S_t)_{t\geq 0}$ and a bank account that offers a constant risk-free rate r. If you were to invest 1 unit in the bank account at time 0, the value of that unit at time t is given by $N_t = e^{rt}$. We also assume that under the real world measure, $(S_t)_{0\leq t\leq T}$ satisfies the Itô process adapted to the filtration $(\mathcal{F}_t)_{0\leq t\leq T}$

$$dS_t = \mu S_t dt + \sigma_t S_t dB_t \quad S_0 > 0$$

where B is a Brownian motion, μ is a constant and σ_t is a non-anticipative processes. We will focus on non path-dependent European contingent claims here, however the main ideas generalise to more complex products. Suppose we have a contingent claim whose payoff $V(S_T)$ is a function of the terminal value of the asset.

A *risk neutral pricing measure*, is an **equivalent** measure such that the discounted asset $e^{-rt}S_t$ is a martingale. Here, equivalent means that the null sets agree i.e. two measures $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(2)}$ on a measurable space $(\Omega, \mathcal{B})$ are equivalent if

$$\mathbb{P}^{(1)}(A) = 0 \quad \Leftrightarrow \quad \mathbb{P}^{(2)}(A) = 0 \qquad \text{for all } A \in \mathcal{B}.$$

This is the first key point, in classical derivative pricing one usually focuses on equivalent measures to the 'real world' measure. In fact we do not know precisely what this 'real world' measure is, we usually assume some model like a geometric Brownian motion, or some stochastic volatility model. However this is quite restrictive, since, making an assumption like this can cause your final results to differ from what is observed in the market. For instance, the assumption of a geometric Brownian motion as in the Black-Scholes model yields a flat volatility surface, which is not the case in reality. Nevertheless, for now we focus only on equivalent measures; the first fundamental theorem of asset pricing is as follows

Theorem 1.1.1 (First Fundamental Theorem of Asset Pricing)**.** *The market is arbitrage-free if and only if there exists a risk neutral pricing measure.*

Thus if we introduce another asset V_t whose terminal value is $V(S_T)$, to maintain an arbitrage-free market there must still exist a risk neutral pricing measure such that $e^{-rt}V_t$ is also a martingale process. Suppose this holds true, then under this risk neutral pricing measure, the martingale property suggests

$$e^{-rt}V_t = \mathbb{E}\left[e^{-rT}V(S_T)|\mathcal{F}_t\right] \quad \Rightarrow \quad V_t = \mathbb{E}\left[e^{-r(T-t)}V(S_T)\Big|\mathcal{F}_t\right].$$

This is known as the *risk neutral pricing formula* which provides an arbitrage-free price for the derivative V. Notice that there may be multiple risk neutral pricing measures, thus we may have multiple arbitrage-free prices.

Often, an institution may want to hedge the derivative so as to cover their losses from the obligations of the contract. In other words, we must find a self-financing replicating strategy (adapted) $(h_t)_{0\leq t\leq T}$ for the derivative V such that

$$V_T - h_0 S_0 - \int_0^T h_t dS_t = 0 \qquad \mathbb{P}-a.s,$$

this is known as a *perfect hedge*. Note that the cost of this replicating strategy is $V_0 = h_0 S_0$ i.e. the initial capital required for the self-financing strategy. So we also have $V_T = h_T S_T$. A *complete* market is one such that every contingent claim is replicable. This leads us to the second fundamental theorem of asset pricing,

Theorem 1.1.2 (Second Fundamental Theorem of Asset Pricing)**.** *An arbitrage-free market is complete if and only if there exists a* ***unique*** *risk neutral pricing measure.*

We may find ourselves in the case where there are multiple risk neutral pricing measures, in which case we cannot guarantee the existence of a replicating strategy. Instead we consider the notion of *super-replicating strategies*, which are pairs $(h_t)_{0\leq t\leq T} \in L^2(\Omega)$ (trading strategy) and $a \in \mathbb{R}$ (initial capital) such that

$$a + \int_0^T h_t dS_t \geq V_T \qquad \mathbb{P}-a.s.$$

In other words, when we hedge with a super-replicating strategy we make almost sure profit. Note that since we are only considering equivalent measures, the notion of a super-replicating strategy holds for all of our measures. In other words, if two measures $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(2)}$ are equivalent, and a property P holds $\mathbb{P}^{(1)}-a.s$, then the property P holds $\mathbb{P}^{(2)}-a.s$. When we discuss super-replicating strategies in a more general setting where the pricing measures are no longer equivalent, we have to be more careful about how we define a super-replicating strategy.

From a practical perspective, it is not feasible to perform dynamic trading strategies as defined in the above. In reality we can only re-balance our portfolio at discrete times, furthermore, re-balancing our portfolio may incur transaction costs. Due to these reasons and more, our final value of the replicating strategy is likely to be less than its theoretical value. Thus considering super-replicating strategies, such that the almost sure profits cover the losses from market imperfections is of practical importance.

Suppose that the value of the derivative V_t can be written as a deterministic function f of time to maturity $T-t=\tau$, the asset price $S_t = s$, the instantaneous volatility $\sigma_t = \sigma$ and the risk-free rate $r_t = r$. So we have the function $f(\tau, s, \sigma, r) = V_t$. We summarise these parameters by $\boldsymbol{\theta} = (\tau, s, \sigma, r)$, then for a small perturbation $\Delta\boldsymbol{\theta} = (\Delta\tau, \Delta s, \Delta\sigma, \Delta r)$, informally we perform a Taylor expansion up to the second order to obtain

$$f(\boldsymbol{\theta}+\Delta\boldsymbol{\theta}) - f(\boldsymbol{\theta}) \approx \sum_{1\leq i\leq 4} \frac{\partial f}{\partial \theta_i}(\boldsymbol{\theta})\Delta\theta_i + \sum_{1\leq i<j\leq 4} \frac{1}{2}\frac{\partial^2 f}{\partial \theta_i^2}(\boldsymbol{\theta})\Delta\theta_i\Delta\theta_j.$$

In other words, the PnL (LHS) from a change in the underlying asset S_t can be decomposed into a PnL attributed to each of the *sensitivities* (RHS) known as the **Greeks**. They describe how the value of the derivative V_t changes as the underlying variables change. By considering the payoff structure of a derivative (S_T vs $V(S_T)$), and noticing that the price structure (S_t vs V_t) is almost (informally) a 'smoothing' of the payoff, it can become quite intuitive to estimate how the Greeks behave. Sticking to the above notation, some of the common Greeks are

- **Delta** = $\frac{\partial f}{\partial s}$ describes how the value of the derivative changes with respect to a change in the spot S_t. Geometrically, with respect to the price structure of the derivative, the Delta is the gradient of the curve. Notice for a European call option, the Delta is always between 0 and 1.
- **Gamma** = $\frac{\partial^2 f}{\partial s^2}$ describes how the Delta changes with respect to a change in the spot S_t. Geometrically, this is the convexity of the price structure. Notice that European call and put options have a convex payoff, so the Gamma is always non-negative.
- **Theta** = $-\frac{\partial f}{\partial \tau}$ describes how the value of the derivative changes with respect to time to maturity. Geometrically, the price structure (S_t vs V_t) converges to the payoff structure (S_T vs $V(S_T)$). The sign of Theta is given by the direction in which the price structure converges to the payoff structure as $t \to T$. For European calls and puts, the price converges downward onto the payoff, so Theta is negative.
- **Vega** = $\frac{\partial f}{\partial \sigma}$ describes how the value of the derivative changes with respect to a change in the volatility σ_t. The sign of Vega is often similar to the sign of Gamma, however there are instances where this is not the case.

Understanding how the Greeks evolve (especially Gamma) is important for understanding the Uncertain Volatility Model.

For the rest of this chapter, we introduce the uncertain volatility setting and discuss how this defines a family of **non-equivalent** pricing measures on the contrary to our discussions in these preliminaries. We then study how to find the maximal arbitrage-free price *theoretically* which is intrinsically linked to the notion of super-replicating strategies, which is discussed further in Chapter 2.

1.2 Uncertain Volatility Model

For the rest of this chapter we follow the derivations provided by Avellaneda et al [1]. Let $T > 0$ and $(B_t)_{0 \le t \le T}$ be a Browian motion with $(\mathcal{F}_t)_{0 \le t \le T}$ it natural filtration. Suppose we have a single asset given by the Itô process,

$$dS_t = \mu_t S_t dt + \sigma_t S_t dB_t, \quad S_0 > 0,$$

where μ_t and σ_t are non-anticipative functions. Furthermore, let $0 < \underline{\sigma} \le \overline{\sigma} < \infty$, and impose that,

$$\underline{\sigma} \le \sigma_t \le \overline{\sigma} \quad \text{for all } t \in [0, T]. \tag{1.2.1}$$

Here we have assumed nothing about our volatility process σ_t other than it being bounded and non-anticipative. We will often refer to the process $(\sigma_t)_{0 \le t \le T}$ being *admissible* if it satisfies these conditions.

If we consider a space of paths Ω where each path is defined on $[0, T]$, takes value in $\mathbb{R}$, and starts at S_0. Then an Itô process, as above may be used to define a probability measure on this space of paths. It is important to realise that the above setup actually provides a family of mutually singular probability measures. Take the two functions,

$$\sigma_t^{(1)} = \underline{\sigma} \quad \text{and} \quad \sigma_t^{(2)} = \overline{\sigma} \quad \text{for all } t \in [0, T].$$

The admissible processes given by $\sigma^{(1)}$ and $\sigma^{(2)}$ define non-equivalent probability measures $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(2)}$. In other words, the set of paths that almost surely occur in one measure is a null set with respect to the other measure.

Now suppose in the simplest case we have a derivative whose payoff at maturity T is given by the function $V(S_T)$. Assuming there is no arbitrage, under the respective risk-neutral measure for each of the admissible processes $(\sigma_t)_{0 \le t \le T}$ the asset satisfies

$$dS_t = rS_t dt + \sigma_t S_t dB_t, \quad S_0 > 0,$$

where r is the risk-free rate. In particular this forms a family of non-equivalent pricing measures defined by the family of admissible processes $(\sigma_t)_{0\leq t\leq T}$. Note that although the exact risk neutral measure here depends on the choice of σ_t, the risk-free rate should be the same amongst all of these measures as given by the risk-free bank account. The question we wish to solve now, is given this family of Itô processes, what is the maximal arbitrage-free price.

1.3 Black-Scholes-Barenblatt PDE

With a slight abuse of notation, for a given admissible process $(\sigma_t)_{0\leq t\leq T}$ an arbitrage-free price for the derivative in question is given by the risk-neutral pricing formula

$$V(t,S_t)=\mathbb{E}\left[e^{-r(T-t)}V(S_T)\middle|\mathcal{F}_t\right],$$

where the expectation is taken under the corresponding risk-neutral measure. Thus to find the maximal arbitrage-free price we consider

$$W(t,S_t)=\sup_{(\sigma_s,s\in[t,T])}\mathbb{E}\left[e^{-r(T-t)}V(S_T)|S_t\right], \tag{1.3.1}$$

where the supremum is taken over all admissible processes $(\sigma_t)_{0\leq t\leq T}$. If we treat this as a stochastic control problem, then the process σ_t is our control, and (1.3.1) is the dynamic value function. Then it follows that the function $(t,s)\mapsto W(t,s)$ solves the PDE

$$\sup_{\sigma\in[\underline{\sigma},\overline{\sigma}]}\left[\frac{\partial W}{\partial t}+r\left(s\frac{\partial W}{\partial s}-W\right)+\frac{1}{2}\sigma^2s^2\frac{\partial^2W}{\partial s^2}\right]=0$$
$$W(T,s)=V(s).$$

By solving the maximisation, the optimal control is given by

$$\sigma(t,s)=\begin{cases}\overline{\sigma} & \text{if } \frac{\partial^2W}{\partial s^2}(t,s)\geq 0\\ \\ \underline{\sigma} & \text{if } \frac{\partial^2W}{\partial s^2}(t,s)<0\end{cases} \tag{1.3.2}$$

Thus the maximal arbitrage-free price under uncertain volatility is given by $W(t,S_t)$ where $(t,s)\mapsto W(t,s)$ is the solution to the second order non-linear partial differential equation, dubbed the Black-Scholes-Barenblatt (BSB) PDE

$$\frac{\partial W}{\partial t}+r\left(s\frac{\partial W}{\partial s}-W\right)+\frac{1}{2}\sigma^2(t,s)s^2\frac{\partial^2W}{\partial s^2}=0 \tag{1.3.3}$$
$$W(T,s)=V(s), \tag{1.3.4}$$

where $\sigma(t,s)$ is given by (1.3.2).

Remark 1.3.1. The optimal control here is intuitive, suppose we are holding a European option at time $t<T$ and our gamma is currently positive. In other words the value of the option at time t with the current spot S_t is locally convex. Thus it is clear that to increase the value of our option volatility should be higher.

Remark 1.3.2. The optimal control σ_t in this simple case is a bang-bang function i.e. it only takes values $\underline{\sigma}$ or $\overline{\sigma}$. This is not the case as multiple underlying assets are introduced, however within the contents of this thesis we only consider derivatives based on a single underlying (but with discretely updating path-dependencies).

Chapter 2

Super-Replicating with Uncertain Volatility

In this chapter we aim capture the main ideas presented by Denis and Martini [3]. By allowing the underlying volatility process to remain unspecified, we define a family of possible processes $(S_t)_{0\leq t\leq T}$. A natural question is, given any realisation of these processes, how can we hedge a derivative in such a way that our profits always cover our losses. This is an important goal for financial institutions that sell derivatives and is known as a super-replicating problem. It turns out that the cheapest super-replicating price in certain settings will equal the maximal arbitrage-free price, which our numerical methods are centred around.

One approach to this super-replicating problem is to specify some model, for instance, suppose there is a single asset $(S_t)_{0\leq t\leq T}$ and a risk-free bank account that offers a risk-free rate of r. Furthermore, we assume $(S_t)_{0\leq t\leq T}$ is a geometric Brownian motion for some fixed $\mu \in \mathbb{R}$ and $\sigma > 0$,

$$dS_t = \mu S_t dt + \sigma S_t dB_t, \quad S_0 > 0.$$

It is well known that under this setting there exists a unique measure such that the discounted process $e^{-rt}S_t$ is a martingale. The first fundamental theorem of asset pricing ensures that if there exists a risk neutral pricing measure, then there exists an arbitrage-free price for a contingent claim V_t. Since the measure is unique, we have a unique arbitrage-free price. Furthermore, from the second fundamental theorem of asset pricing, there exists a replicating strategy for the contingent claim.

We could have specified a model for our market such that there are multiple risk neutral pricing measures, so let's instead consider this more general setting. We denote this set of **equivalent** risk neutral pricing measures on the measurable space $(\Omega, \mathcal{B})$ by $\mathbf{P}'_m$. The second fundamental theorem of asset pricing tells us that in this setting the market is incomplete. Since there are multiple pricing measures, we have multiple arbitrage-free prices for a contingent claim V_t. Thus it is now ambiguous which of these prices we should sell at. Our original goal is to hedge the contingent claim, which requires a replicating strategy, however when the market is incomplete there may not exist a replicating strategy and a perfect hedge.

A super-replicating strategy is a pair $a \in \mathbb{R}$ and an adapted process $(h_t)_{0\leq t\leq T} \in L^2(\Omega)$ such that

$$a + \int_0^T h_t dS_t \geq V_T \quad \mathbb{P} - a.s.$$

Notice that by hedging with a super-replicating strategy we $\mathbb{P} - a.s$ make a profit. The value $a \in \mathbb{R}$ is the price of the super-replicating strategy. For bounded payoffs V_T, there always exists a super-replicating strategy. We define the cheapest super-replicating price by the function

$$\Lambda(V_T) = \inf\left\{ a \in \mathbb{R} \middle| \exists h \in L^2(\Omega) \text{ such that } a + \int_0^T h_t dS_t \geq V_T \quad \mathbb{P} - a.s \right\}.$$

There is an important result which provides a link between the set of arbitrage-free prices and cheapest super-replicating price in the case where we consider a family of equivalent measures, the duality formula is given by

$$\Lambda(V_T) = \sup_{\mathbb{P}\in\mathbf{P}'_m} \mathbb{E}^{\mathbb{P}}(V_T). \tag{2.0.1}$$

So in certain settings the cheapest super-replicating price is equal to the largest arbitrage-free price [2, 4, 5, 9].

In our setting where the volatility process $(\sigma_t)_{0\leq t\leq T}$ is unspecified we naturally have multiple pricing measures. The difference here is that these measures are in general **not** equivalent. Consider two such pricing measures $\mathbb{P}_1$ and $\mathbb{P}_2$ that are not equivalent, then suppose we have a super-replicating strategy (a, h) such that

$$a + \int_0^T h_t dS_t \geq V_T \quad \mathbb{P}_1 - a.s.$$

Since the two measures are not equivalent, there exists a set $A \in \mathcal{B}$ such that $\mathbb{P}_1(A) = 0$ but $\mathbb{P}_2(A) > 0$. Then on this set, the strategy (a, h) may not be super-replicating, i.e. for some $B \subseteq A$ with $\mathbb{P}_2(B) > 0$ we have for all $\omega \in B$

$$a + \int_0^T h_t dS_t(\omega) < V_T(\omega).$$

Thus we must now be careful with how we define the notion of $\mathbb{P} - a.s$ profit. The idea, by Denis and Martini, is to build a framework which allows us to determine when events occur almost surely uniformly on sets of probability measures that may not be equivalent. The end goal of this framework, is to develop an equivalent characterisation for the (2.0.1) which will provide the theoretical grounding for our methods for pricing with uncertain volatility. More specifically, in trying to find the cheapest super-replicating price, we will instead consider the simpler problem of finding the maximal arbitrage-free price.

For the following we adopt the notation of a super-replicating strategy being a pair (a, h) such that

$$a + \int_0^T h_t dZ_t \geq V_T \quad \mathbb{P} - a.s,$$

where in the above example where S_t followed a Geometric Brownian Motion,

$$dS_t = S_t dZ_t$$
$$Z_t = \sigma B_t + \mu t.$$

This is equivalent to our above definition where we instead define the strategy $\tilde{h}_t = h_t S_t$. The advantage with this alternative view is that $Z_0 = 0$ which is simply more convenient.

2.1 An Alternative Framework for Model Uncertainty

Let $\Omega = C([0, T], \mathbb{R})$, the space of scalar continuous functions $Z = (Z_t)$ and $Z_0 = 0$, be endowed by the uniform norm and let $\mathcal{B}$ be its Borel σ-field. We denote by $(\mathcal{F}_t)_{t\geq 0}$ the canonical filtration. A probability measure $\mathbb{P}$ on the measurable space $(\Omega, \mathcal{B})$ is a **martingale measure** if the process Z is a martingale with respect to $(\mathcal{F}_t)_{t\geq 0}$ under $\mathbb{P}$. We denote by $\mathbf{P}_m$ the set of all such martingale measures.

This setting so far is fairly general, in the bigger picture we look at all possible continuous paths that start at 0. We then look at measures that assign positive probability to sets of paths where under this same measure the path obeys the martingale property. We may have two measures $\mathbb{P}_1, \mathbb{P}_2 \in \mathbf{P}_m$ and two disjoint sets $A_1, A_2 \in \Omega$ such that $\mathbb{P}_1(A_1) = \mathbb{P}_2(A_2) = 1$ but $\mathbb{P}_1(A_2) = \mathbb{P}_2(A_1) = 0$, thus these two measures are not equivalent. In this case there is no change of measure that can take us from $\mathbb{P}_1$ to $\mathbb{P}_2$ which makes it difficult to work with the typical set-up of stochastic integration. We denote the quadratic variation of Z under $\mathbb{P}$ by $\langle Z\rangle^{\mathbb{P}}$ which is defined up to $\mathbb{P}$-null sets. Now

we fix two nonzero measures $\overline{\mu}, \underline{\mu}$ on $[0,T]$. We denote the continuous distribution function by $\mu_t = \mu([0,t])$. Then we consider the subset $\mathbf{P} \subseteq \mathbf{P}_m$ of martingale measures that satisfy the following assumption.

Assumption 2.1.1 ($H(\underline{\mu}, \overline{\mu})$)**.** For each $\mathbb{P} \in \mathbf{P}$, the processes $\langle Z \rangle^{\mathbb{P}} - \underline{\mu}$ and $\overline{\mu} - \langle Z \rangle^{\mathbb{P}}$ are increasing up to $\mathbb{P}$-null sets. This relationship is denoted by

$$d\underline{\mu}_t \leq d\langle Z \rangle_t^{\mathbb{P}} \leq d\overline{\mu}_t$$

It is in this context that we are applying bounds to the volatility of our underlying asset. In our particular case we consider $d\underline{\mu}_t = \underline{\sigma}^2 dt$ and $d\overline{\mu}_t = \overline{\sigma}^2 dt$. So this framework is a generalisation of the original setting introduced by Avellaneda et al [1] and Lyons [11].

Notice the payoff of a contingent claim V_T can be written as a function $f \in C_b(\Omega)$ — the set of continuous bounded function on Ω. The framework, as presented by Denis and Martini, is built on the notion of a *capacity*, beyond the definition and some basic vocabulary we do not want to delve too far into the technicalities as the overall goal here is to illustrate the fundamental difficulties in expressing a duality formula (2.0.1) when there is model uncertainty.

For each $f \in C_b(\Omega)$, we define the **capacity**, a convex positive homogeneous function $c : C_b(\Omega) \to \mathbb{R}$ by

$$c(f) = \sup\left\{ ||f||_{L^2(\Omega,\mathbb{P})} \mid \mathbb{P} \in \mathbf{P} \right\} = \sup_{\mathbb{P}\in\mathbf{P}} \left(\int_\Omega f(Z)^2 d\mathbb{P}(Z) \right)^{\frac{1}{2}}.$$

Furthermore, for some $A \in \Omega$ we set $c(A) = c(\mathbb{1}_A)$. Then a set A is **polar** if $c(A) = 0$ and say a property holds "quasi-surely" (q.s.) if it holds outside a polar set. Notice that if A is measurable and polar, then $\mathbb{P}(A) = 0$ for all $\mathbb{P} \in \mathbf{P}$. Thus to say a property holds q.s. is to say the property holds a.s. for all measures $\mathbb{P} \in \mathbf{P}$. To make it clear why this is a useful notion for us, one of the previously highlighted problems when considering a super-replicating strategy (a, h) for non-equivalent measures is that for one measure the strategy may be super-replicating however this may not hold for other measures. Instead if we can find a super-replicating strategy (a, h) such that

$$a + \int_0^T h_t dZ_t \geq V_T \quad \text{q.s.} \tag{2.1.1}$$

Then there is no question of whether hedging with this strategy will lead to certain profits as it does under all valid measures $\mathbb{P} \in \mathbf{P}$.

There is a technicality here that still needs to be taken into consideration, when defining a super-replicating strategy (a, h), the process h_t must be admissible integrand to ensure the stochastic integral is defined up to $\mathbb{P}$-null sets. The problem is that for two measures $\mathbb{P}_1, \mathbb{P}_2 \in \mathbf{P}$, the process h_t may be an admissible integrand to integrate under the measure $\mathbb{P}_1$, however, it may not be admissible when integrating under the measure $\mathbb{P}_2$. Thus, the space of admissible integrands needs to be chosen with care, much of the work in developing this framework goes into re-developing a theory of stochastic integration under a set of non-equivalent measures $\mathbf{P}$.

2.2 Stochastic Integration for Model Uncertainty

The following is analogous to developing a theory of stochastic integration for a given probability measure $\mathbb{P}$, the goal of this section is to highlight the differences required when we want to allow integration across a set of non-equivalent measures $\mathbf{P}$. Classically you might define the integral over a suitable set of integrands for a given measure $\mathbb{P}$, then one can extend this integral over equivalent probability measures $\mathbb{Q}$ via the Radon-Nikodym derivative. As for reasons previously mentioned, it is no longer trivial how we can extend this integral when measures are not equivalent.

Notice that the capacity c defines a semi-norm on the space of bounded functions $C_b(\Omega)$, we denote by $\mathcal{L}$ the topological completion of $C_b(\Omega)$ under this semi-norm, furthermore we denote by L the quotient of $\mathcal{L}$ with respect to the quasi-sure equivalence relation. There is a similarity here to the

construction of Lebesgue measurable functions, instead of choosing a single measure λ and defining the space of integrable functions under this measure, we instead take all of the measures $\mathbb{P} \in \mathbf{P}$ and define the space of functions that are integrable under all of the measures $\mathbb{P} \in \mathbf{P}$.

Theorem 2.2.1 ([3, Theorem 2.1]). *Each element in $\mathcal{L}$ can be identified with a quasi-continuous function on Ω (and so is defined quasi-surely). Moreover, (L, c) is a Banach space.*

Notice that from the definition of the capacity, we have for all $\mathbb{P} \in \mathbf{P}$

$$||f||_{L^2(\Omega,\mathbb{P})} \leq c(f) \quad \text{for } f \in C_b(\Omega).$$

If $f \in \mathcal{L}$, then by definition there is a sequence $f_n \in C_b(\Omega)$ that converges to f in $\mathcal{L}$. So, f_n is also a Cauchy sequence in $L^2(\Omega, \mathbb{P})$ for each $\mathbb{P} \in \mathbf{P}$ and thus converges to some function in $L^2(\Omega, \mathbb{P})$ equal to f $\mathbb{P}$-a.s. Consequently, we can extend the domain of the capacity c to $\mathcal{L}$

Proposition 2.2.2 ([3, Proposition 2.2]). *Let $f \in \mathcal{L}$. Then $c(f) = \sup\left\{||f||_{L^2(\Omega,\mathbb{P})} \mid \mathbb{P} \in \boldsymbol{P}\right\}$.*

Looking at the bigger picture, the set of functions $\mathcal{L}$ defines a large class of payoffs as functions of paths Ω. It should be clear that building a good understanding of this space in the case where we have a large set of non-equivalent measures $\mathbb{P}$ is necessary. The next step is to understand how trading strategies h fit into the picture, in particular we want to formulate a notion of integrating these strategies against paths in Ω to super-replicate payoffs in $\mathcal{L}$.

Let $\mathcal{H}_e$ be the set of "elementary" processes $h_s = \sum_{i=0}^{N-1} k_{t_i} \mathbb{1}_{[t_i,t_{i+1})}(s)$, where $0 = t_0 < t_1 < ... < t_N = T$ is a partition of $[0, T]$ and k_{t_i} are $\mathcal{F}_{t_i}$-measurable, bounded and continuous. Then by $\mathcal{H}$ we denote the completion of $\mathcal{H}_e$ with respect to the semi-norm

$$||h||_{\mathcal{H}} = c\left(\left(\int_0^T h_s^2 d\overline{\mu}_s\right)^{\frac{1}{2}}\right) = \sup_{\mathbb{P}\in\mathbf{P}} \mathbb{E}^{\mathbb{P}}\left(\int_0^T h_s^2 d\overline{\mu}_s\right)^{\frac{1}{2}},$$

and denote by H the quotient of $\mathcal{H}$ with respect to the linear space of processes such that $||h||_{\mathcal{H}} = 0$. Notice that the space $\mathcal{H}_e$ is just the set of trading strategies with a finite number of portfolio re-balances. By taking the completion $\mathcal{H}$, we are considering more general dynamic trading strategies.

The idea here is that if we can define the integral against $\overline{\mu}_t$ over this class of functions $\mathcal{H}$, then since the assumption $H(\underline{\mu}, \overline{\mu})$ holds we can integrate against the process Z_t. Moreover, the function $h \mapsto \int_0^T h_t dZ_t$ is contained in $\mathcal{L}$. In otherwords, the strategy h provides a quasi-sure replicating strategy for some payoff f. Full details are given in Section 2.2.1, [3], here we state the theorem that defines the stochastic integral

Theorem 2.2.3 (Theorem 2.8, [3]). *The linear mapping*

$$h = \sum_{i=0}^{N-1} k_{t_i} \mathbb{1}_{[t_i,t_{i+1})} \mapsto I_T(h) = \int_0^T h_s dZ_s = \sum_{i=0}^{N-1} k_{t_i}\left(Z_{t_{i+1}} - Z_{t_i}\right),$$

considered as a function from $\mathcal{H}_e$ to $\mathcal{L}$, admits the bound

$$c\left(I_T(h)\right) \leq ||h||_{\mathcal{H}} \tag{$*$}$$

It can be extended uniquely to a continuous linear mapping from $\mathcal{H}$ to $\mathcal{L}$, still denoted by $I_T(h) = \int_0^T h_s dZ_s$, that satisfies ($*$).

These spaces of functions form a coherent framework to tackle the problem of super-replicating strategies among sets of non-equivalent pricing measures. In short, the idea is similar to developing a theory of stochastic integration for a single measure, however, instead we limit the set of admissible functions to sets of functions that are admissible when integrating against all of the measures that satisfy the assumption $H(\underline{\mu}, \overline{\mu})$.

2.3 The Duality Formula

Recall that a function $f \in \mathcal{L}$ can be interpreted as a contingent claim. Following Denis and Martini (Section 2.3, [3]), we define $K = \{I_T(h) | h \in \mathcal{H}\}$, the terminal values of portfolio processes h. Then in our context of model uncertainty, we can define the cheapest super-replicating price as

$$\Lambda(f) = \inf \{a \in \mathbb{R} | \exists g \in K \text{ such that } a + g \geq f \quad q.s\} .$$

Then the duality formula can be shown to hold for a large class of payoffs $\Gamma \subseteq C_b(\Omega)$.

Theorem 2.3.1 ([3, Theorem 6.1]). *Let $\underline{\mu}$ and $\overline{\mu}$ be two deterministic measures on $[0,T]$ such that $d\underline{\mu} \leq d\overline{\mu}$ and $\overline{\mu}$ is Hölder continuous. Let $\boldsymbol{P}$ be the set of all martingale measures which satisfy the assumption $H(\underline{\mu}, \overline{\mu})$. Let f be a bounded continuous function in Γ. Then*

$$\Lambda(f) = \sup \left\{ \mathbb{E}^{\mathbb{P}} f \middle| \mathbb{P} \in \boldsymbol{P} \right\} .$$

To understand some of the functions that are contained in the class of functions Γ, we have the following lemmas,

Lemma 2.3.2 ([3, Lemma 5.4]). *If f is a continuous and bounded cylindrical function, then $f \in \Gamma$.*

Lemma 2.3.3 ([3, Lemma 5.5]). *Let $F : \mathbb{R} \to \mathbb{R}$ be a continuous function and let $G : \mathbb{R} \to \mathbb{R}$ be a bounded continuous function. Then $f = G\left(\int_0^T F(Z_s) ds\right)$ is in Γ.*

Lemma 2.3.4 ([3, Lemma 5.6]). *Let $S = \sup_{t \in [0,T]} Z_t$ and $G : \mathbb{R} \to \mathbb{R}$ be a bounded continuous function. Then $f = G(S)$ is in Γ.*

In short, the duality formula in Theorem 2.3.1 covers a large class of general path-dependent European payoffs. This is an important result — although a large focus on uncertainty in volatility has been to find the maximum arbitrage-free price, it is important to remember that the original goal is to hedge the derivative; for this we require a super-replicating strategy. Practically, obtaining the price of the cheapest super-replicating strategy is challenging, however obtaining the maximal arbitrage-free price is a more manageable task. To this end, we continue by investigating different methods for obtaining this maximal arbitrage-free price and thus the cheapest super-replicating price.

Chapter 3

Zero Gamma Boundaries

The work presented in this chapter is my own. In the single asset case it is clear that the regions in $[0,T] \times \mathbb{R}^+$ where gamma is positive or negative is important for determining where in our simulation we should choose volatility of $\overline{\sigma}$ or $\underline{\sigma}$. For our studies we consider $z : [0,T] \to \mathbb{R}^+$ such that the gamma is zero along these boundaries, as these will help us characterise these regions. Some care is required when we talk about the UVM 'gamma', since uncertain volatility does not specify an exact model, but rather a family of models.

We are making an assumption here that there are continuous paths z that separate the regions into positive and negative gamma, however for the cases we care about this will be fairly evident. For the following we will focus on non path-dependent options, the results that we arrive at will still hold in cases where the path dependencies update at deterministic discrete times (since close to expiry this is the same case by case as a non path-dependent option). For the following we make a simplifying assumption that the risk-free rate is 0.

Definition 3.0.1. For an asset $(S_t)_{0 \le t \le T}$ and a derivative with payoff $V(S_T)$. Assuming no arbitrage, there exists a price for the derivative at time $t \in [0,T]$ as a function of S_t denoted by $V(t, S_t)$. A **Zero Gamma Boundary** is a continuous path $z : [0,T] \to \mathbb{R}^+$ that satisfies

$$\frac{\partial^2 V}{\partial s^2}(t, z(t)) = 0 \quad \text{for all } t \in [0,T].$$

Moreover, in the Uncertain Volatility setting, a Zero Gamma Boundary is a continuous path $z : [0,T] \to \mathbb{R}^+$ that satisfies

$$\frac{\partial^2 W}{\partial s^2}(t, z(t)) = 0 \quad \text{for all } t \in [0,T].$$

Where $(t,s) \mapsto W(t,s)$ is the solution to the Black-Scholes-Barenblatt equation (1.3.3) with terminal condition $W(T,s) = V(s)$.

Our goal is to show that the terminal values of these zero gamma boundaries, $z(T)$, agree regardless of which model we choose.

Remark 3.0.2. For a given derivative there may be multiple zero gamma boundaries, or they may simply not exist as with a vanilla call option. Furthermore, these paths depend on the underlying choice of model. For a call spread with strikes K_d and K_u where the underlying model is Black-Scholes with volatility σ and risk-free rate $r = 0$, there is a unique zero gamma boundary given by

$$z(t) = \sqrt{K_d K_u} e^{-\frac{\sigma^2}{2}(T-t)} \quad \text{for } t \in [0,T].$$

This is a straight forward derivation from the Black-Scholes formula for a European call option. Notice that the endpoint $z(T)$ is equal to the geometric mean of the two strikes which is independent of the free parameters to choose from in the Black-Scholes model.

Conjecture 3.0.3. For a given payoff $V(S_T)$, the terminal values of the zero gamma boundaries, $z(T)$, are independent of the model choice for $(S_t)_{0 \le t \le T}$.

We show that in a Black-Scholes setting for all derivatives with payoff $V(S_T)$, the terminal value of a zero gamma boundary $z(T)$ is independent of σ and thus agrees across all Black-Scholes models. This is a good initial check for the validity of Conjecture 3.0.3. The core idea of the proof is to show that by adjusting the volatility σ, the zero gamma boundary is simply adjusted by scaling the time variable.

Lemma 3.0.4. *Let $z : [0,T] \to \mathbb{R}^+$ be a zero gamma boundary for some derivative with payoff $V(S_T)$ where the underlying model for pricing is Black-Scholes with volatility $\sigma > 0$. Then $z(T)$ is independent of σ.*

Proof. Let $\sigma^{(2)} > \sigma^{(1)} > 0$ and define $\xi = \left(\frac{\sigma^{(2)}}{\sigma^{(1)}}\right)^2 > 1$. Suppose we have two processes satisfying

$$\begin{aligned} dS_t^{(1)} &= \sigma^{(1)} S_t^{(1)} dB_t^{(1)}, \quad t \in [0, \xi T] \\ dS_t^{(2)} &= \sigma^{(2)} S_t^{(2)} dB_t^{(2)}, \quad t \in [0, T] \\ S_0^{(1)} &= S_0^{(2)} = S_0 \in (0, \infty), \end{aligned}$$

where $B^{(1)}$ and $B^{(2)}$ are independent Brownian motions. Suppose we have a payoff $V(\cdot)$. Then conditional on $S_t^{(i)} = u$, let $V^{(1)}(t,u)$ be the fair price for the derivative whose payoff is $V(S_{\xi T}^{(1)})$ and let $V^{(2)}(t,u)$ be the fair price for the derivative whose payoff is $V(S_T^{(2)})$. Notice that the two derivatives has different maturities $\xi T > T$.
Fix $\Delta t \in [0, \xi T]$, set $\Delta t' = \Delta t / \xi$ then we have the equalities in distribution

$$\begin{aligned} S_{\xi T}^{(1)} &= S_{\xi T - \Delta t}^{(1)} \exp\left(\sigma^{(1)} \sqrt{\Delta t} Z^{(1)} - \frac{1}{2}\left(\sigma^{(1)}\right)^2 \Delta t\right), \\ S_T^{(2)} &= S_{T - \Delta t'}^{(2)} \exp\left(\sigma^{(2)} \sqrt{\Delta t'} Z^{(2)} - \frac{1}{2}\left(\sigma^{(2)}\right)^2 \Delta t'\right), \end{aligned}$$

where $Z^{(1)}$ and $Z^{(2)}$ are independent standard normal random variables. Then conditional on $S_{\xi T - \Delta t}^{(1)} = S_{T - \Delta t'}^{(2)} = u$, we have $S_{\xi T}^{(1)}$ and $S_T^{(2)}$ are equal in distribution. Thus we have the equality in prices

$$V^{(1)}\left(\xi T - \Delta t, u\right) = V^{(2)}\left(T - \frac{1}{\xi}\Delta t, u\right) \quad \text{for all } u \in \mathbb{R}^+.$$

In particular, by setting $t = T - \frac{1}{\xi}\Delta t \in [0, T]$

$$V^{(1)}\left(\xi t, u\right) = V^{(2)}\left(t, u\right) \quad \text{for all } u \in \mathbb{R}^+.$$

Since the equality holds for all u, we also have the equality of gammas

$$\frac{\partial^2 V^{(1)}}{\partial u^2}\left(\xi t, u\right) = \frac{\partial^2 V^{(2)}}{\partial u^2}\left(t, u\right). \qquad (*)$$

Let $z : [0, \xi T] \to \mathbb{R}^+$ be a zero gamma boundary for $S^{(1)}$ i.e.

$$\frac{\partial^2 V^{(1)}}{\partial u^2}(t, z(t)) = 0 \quad \text{for all } t \in [0, \xi T].$$

The goal is to now construct an analogous zero gamma boundary for $S^{(2)}$, and show that the terminal values are equal. Define the path $\tilde{z} : [0,T] \to \mathbb{R}^+$ by

$$\tilde{z}(t) = z\left(\xi t\right) \quad \text{for } t \in [0, T]. \qquad (**)$$

Then it follows

$$\begin{aligned} \frac{\partial^2 V^{(2)}}{\partial u^2}(t, \tilde{z}(t)) &= \frac{\partial^2 V^{(2)}}{\partial u^2}(t, z(\xi t)) \quad \text{by } (**) \\ &= \frac{\partial^2 V^{(1)}}{\partial u^2}(\xi t, z(\xi t)) \quad \text{by } (*) \\ &= 0 \end{aligned}$$

Where the final equality holds since z is gamma neutral for $S^{(1)}$. Since the above holds for all $t \in [0, T]$, we have shown that $\tilde{z}$ is a zero gamma boundary for $S^{(2)}$. In particular we have $\tilde{z}(T) = z(\xi T)$, thus the two paths agree at their terminal values. Since the choice of zero gamma boundary and volatilities $\sigma^{(2)} > \sigma^{(1)} > 0$ were general, the statement must hold for all volatilities and zero gamma boundaries. Thus we conclude that the terminal value of any zero gamma boundary is independent to the choice of σ. □

We now discuss the main ideas that support Conjecture 3.0.3.

- One can imagine that as we approach expiry in the Uncertain Volatility Model, the probability of the volatility changing from $\overline{\sigma}$ to $\underline{\sigma}$ or vice versa goes to zero, thus we could argue that the price at this point is 'ϵ-close' to the price in Black-Scholes. Since we showed the terminal values of the gamma agree in Black-Scholes we could argue that this means they agree with the UVM case — we are only concerned with this case.

- Another idea is, instead of considering the payoff $V(S_T)$, let $\epsilon > 0$ be small, and denote by $V_{BS}(T - \epsilon, s) := V^{\epsilon}_{BS}(s)$ the Black-Scholes price at $T - \epsilon$; close to expiry. Then we instead consider the payoff to be this Black-Scholes price, $V^{\epsilon}_{BS}(S_T)$. For a large class of European options, for example a call spread, there will be finite number of points where the second order derivative with respect to the asset is 0, as a consequence of continuity the zero gamma boundaries must terminate at these points. The payoff $V^{\epsilon}_{BS}(S_T)$ can be made arbitrarily 'close' to the payoff $V(S_T)$.

Chapter 4

Numerical Methods

There are various approaches to computing the Uncertain Volatility price for derivatives presented in literature. Avellaneda at al [1] presented a trinomial tree approach equivalent to an explicit scheme for numerically solving the BSB equation for simple payoffs with no path-dependencies. Such a scheme will be useful to benchmark alternative approaches, however will not be feasible to price more complex products such as cliquets. Guyon and Labordère [8] outline more general approaches that are capable of pricing complex products, even when there are multiple underlying. They present a parametric approach, which in short is a scheme to directly perform the maximisation making use of the dynamic programming principle. The method which we apply to the call spread example and later generalise for cliquets is heavily inspired by this method. They also provide an approach taking advantage of a link between second order Backward Stochastic Differential Equations and second order non-linear parabolic partial differential equations such as the BSB equation.

4.1 Comparative Quantities

We will present and compare multiple methods for varying products, so it is important to construct informative dimensionless quantities in which we can compare methods and the prices they produce. Since the Black-Scholes model with a flat volatility of $\sigma = \frac{\underline{\sigma}+\overline{\sigma}}{2}$ is included in the family of models that the uncertain volatility model considers, the maximal arbitrage-free price will always provide a premium on-top of this Black-Scholes price. Thus in future comparisons, we will often compare prices of products based on the percentage premium from pricing with uncertain volatility

$$\langle \text{Method} \rangle \text{ Premium} = \left(100 \times \frac{\text{UVM}_{\langle \text{Method} \rangle} - \text{BS}_{\sigma}}{\text{BS}_{\sigma}} \right) \%, \qquad \sigma = \frac{\underline{\sigma} + \overline{\sigma}}{2}.$$

This dimensionless variable provides a better insight for comparisons of the price with uncertain volatility across **different** products. Furthermore, to compare the prices produced by various methods, we will use the Trinomial Tree (Section 4.2) as a benchmark, thus we will consider the dimensionless quantity

$$\text{Model Comparison} = \frac{\langle \text{Method} \rangle \text{ Premium}}{\text{Tree Premium}}.$$

So given that the tree method is accurate and should provide an upper bound on the premium, we expect this quantity to normally lie in $[0, 1]$ — 1 being the best result. Note however, due to randomness and inaccuracies from the tree, it is possible that the Method Comparison lies outside of this range for a given model.

4.2 Trinomial Tree

4.2.1 Method

We first present the method given in [1] and outline how the scheme can be adapted to path dependent derivatives. Suppose we have a derivative that has a payoff $V(S_T)$ at some maturity $T > 0$. Furthermore, assume S_t satisfies an SDE given by

$$\frac{dS_t}{S_t} = rdt + \sigma_t dB_t, \quad S_0 > 0, \tag{4.2.1}$$

where σ_t is a non-anticipative process satisfying $\sigma_t \in [\underline{\sigma}, \overline{\sigma}]$.

The following is a discretization of the BSB equation which can be viewed as a trinomial tree that approximates (4.2.1). We begin by discretising the time interval $[0, T]$ into N trading periods that have length $\Delta t = T/N$. Then set

$$\begin{aligned} u &= \exp\left(\overline{\sigma}\sqrt{\Delta t} + r\Delta t\right), \\ m &= \exp\left(r\Delta t\right), \\ d &= \exp\left(-\overline{\sigma}\sqrt{\Delta t} + r\Delta t\right) \end{aligned}$$

We define the possible set of trajectories for the discritised asset price $(S_{t_i})_{i=0}^N$ where after each trading period n the asset price S_{t_n} can transition to one of the three possible values: uS_{t_n}, mS_{t_n} or dS_{t_n}. Note that the values for u, m and d are chosen such that the possible set of trajectories forms a recombining trinomial tree. As a consequence, there are $2n+1$ possible values that S_{t_n} can take for each $n = 0, 1, ..., N$. We can label each node in the tree according to the multi-index (n, j) where $n \in \{0, ..., N\}$ and $j \in \{1, ..., 2n+1\}$. So the first index denotes the trading period and the second denotes indicates the asset price level. Then from node (n, j) we can transition to one of the three nodes (up-middle-down) $(n+1, j+2)$, $(n+1, j+1)$ or $(n+1, j)$. Then the price at node (n, j) is given by

$$S_n^j = S_0 \exp\left(j\overline{\sigma}\sqrt{\Delta t} + nr\Delta t\right).$$

Now define the family of transition probabilities

$$\begin{aligned} P_u(p) &= p\left(1 - \frac{\overline{\sigma}\sqrt{\Delta t}}{2}\right), \\ P_m(p) &= 1 - 2p, \\ P_d(p) &= p\left(1 + \frac{\overline{\sigma}\sqrt{\Delta t}}{2}\right), \end{aligned}$$

for $p \in \left[\frac{\underline{\sigma}^2}{2\overline{\sigma}^2}, \frac{1}{2}\right]$. Then at each step the variance of the log return of the asset $\log \frac{S_{t_n}}{S_{t_{n-1}}}$ is given by

$$\underline{\sigma}^2 \Delta t \leq 2p\overline{\sigma}^2 \Delta t \leq \overline{\sigma}^2 \Delta t$$

Therefore, by controlling the parameter p we can approximate stochastic processes that lie in the family of processes given by (4.2.1). Note that the free parameter p for each transition of node n to $n+1$ may only be chosen at any step $i \leq n$ since the process σ_t is non-anticipative.

Now, in the uncertain volatility setting, we denote the maximal arbitrage-free price at node (n, j) by W_n^j. Then we have

$$W_N^j = V(S_N^j) \quad \text{ for } j \in \{1, ..., 2N+1\}.$$

Notice that since W is the solution to the maximisation problem (1.3.1) we have in the discrete setting

$$\begin{aligned} W_{n-1}^j &= \sup \mathbb{E}\left[e^{-r(N-(n-1))\Delta t} V(S_N) | S_{n-1} = S_{n-1}^j\right] \\ &= \sup \mathbb{E}\left[e^{-r(N-(n-1))\Delta t} V(S_N)\left(\mathbb{1}_{S_n = S_n^{j+2}} + \mathbb{1}_{S_n = S_n^{j+1}} + \mathbb{1}_{S_n = S_n^j}\right) | S_{n-1} = S_{n-1}^j\right] \\ &= e^{-\Delta t} \sup\left(P_u(p) W_n^{j+2} + P_m(p) W_n^{j+1} + P_d(p) W_n^j\right) \end{aligned}$$

where the supremum is taken over the parameters p for each of the trading periods. Notice that the p dependence of each of the transition probabilities is linear, thus we can easily solve the maximisation above. To this end define

$$L_n^j = \left(1 - \frac{\overline{\sigma}\sqrt{\Delta t}}{2}\right) W_n^{j+2} + \left(1 - \frac{\overline{\sigma}\sqrt{\Delta t}}{2}\right) W_n^j - 2W_n^{j+1}$$

Then we have

$$W_{n-1}^j = e^{-r\Delta t} \begin{cases} W_n^{j+1} + \frac{1}{2} L_n^j & \text{if } L_n^j \geq 0 \\ W_n^{j+1} + \frac{\underline{\sigma}^2}{2\overline{\sigma}^2} L_n^j & \text{if } L_n^j < 0 \end{cases}$$

We then apply this relation iteratively until we obtain the desired quantity W_0^0.

4.2.2 Results

We are particularly interested in the call spread example as this will be important for understanding how to price the cliquet. For two strikes $0 < K_d < K_u < \infty$ the payoff of a call spread is given by

$$(S_T - K_d)^+ - (S_T - K_u)^+$$

Furthermore the payoff of the butterfly is given by

$$(S_T - K_d)^+ - 2\left(S_T - \frac{K_d + K_u}{2}\right)^+ + (S_T - K_u)^+$$

We know that as $N \to \infty$ in the trinomial scheme we have convergence of the model to the SDE and thus convergence of the price. Clearly we must settle on some large N, however since this model will be our benchmark it is worth understanding for what range of contingent claims this method of pricing is accurate. For the following define $\Delta K = K_u - K_d$. Now by taking $\overline{\sigma} = \underline{\sigma} = \sigma > 0$ the trinomial scheme is a method for obtaining the Black-Scholes price of a contingent claim. Of which for a call-spread and butterfly we have an analytical formula to compare the scheme against.

In Figure 4.1 we observe the percentage difference between the scheme and the analytical formula for both the call spread and butterfly. As we can see for the call spread the results are extremely accurate, even with $N = 100$, which is fairly small, the percentage difference never deviates more that 0.25% from the Black-Scholes formula. For the butterfly the results are quite extreme, as $\Delta K \to 0$ we observe that the trinomial tree fails to accurately obtain the correct prices. In fact, when $N = 100$ the error is always greater than 1%. The dashed lines on the plot for the butterfly indicate the points where the percentage difference is $\approx 1\%$.

Fortunately, we are more interested in the call spread, of which the prices the trinomial tree produce are accurate. Note that our comparison above is between Black-Scholes prices, however when we introduce uncertainty i.e. $\underline{\sigma} < \overline{\sigma}$ the above errors may change. Nevertheless, the above analysis is certainly an indicator of what errors to expect when pricing in uncertain volatility with the trinomial tree.

In Figure 4.2 we have prices for a particular family of options, outlining the premium offered by pricing with uncertain volatility as oppose to pricing under the Black-Scholes flat volatility of $\frac{\underline{\sigma}+\overline{\sigma}}{2}$.

Unfortunately, this method is only feasible up to options with very few, if any path-dependencies. In particular this drawback motivates the development of Monte Carlo methods as follows.

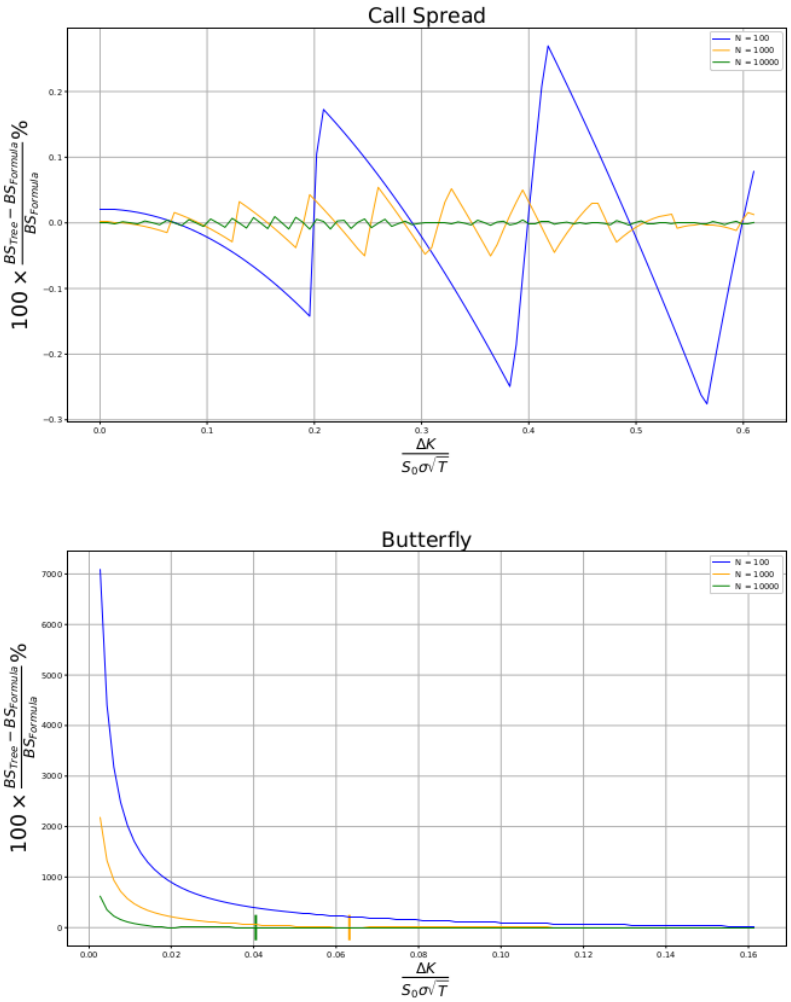

Figure 4.1: Numerical error of the trinomial tree (Black-Scholes) for multiple call spread (top) and butterfly (bottom) options. The y-axis is the percentage difference between the trinomial tree method and Black-Scholes formula. On the x-axis we plot the dimensionless variable $\frac{\Delta K}{S_0\sigma\sqrt{T}}$. On the butterfly plot, the dashed lines represent the first point the percentage difference dips below 1%.

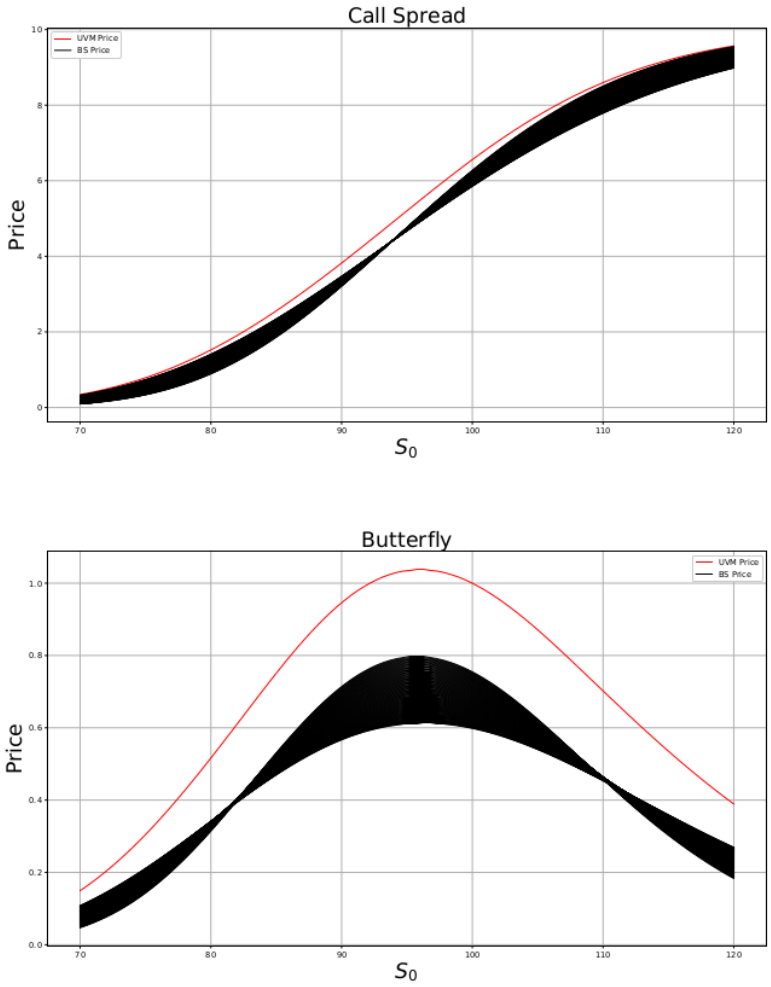

Figure 4.2: A plot of prices where the uncertain volatility is $13\% \leq \sigma_t \leq 17\%$ and Black-Scholes prices with volatility in the same range. Other parameters used are $K_d = 90$, $K_u = 100$, $T = 1$, $r = 0$, $N = 10,000$

4.3 Parametric Approach

We first present the method by Guyon and Labordère [8] and explore an approach for how it can be improved depending on the derivative. In our approach we consider a simple call spread and later generalise this for pricing cliquets. For simplicity, we assume that the risk-free rate $r = 0$.

4.3.1 The Algorithm

The overall idea is instead of dealing with the BSB equation directly, consider the value function (1.3.1) and apply dynamic programming principle in order to optimise directly. More formally we discretize the time interval $[0, T] = \cup_{i=1}^{n}[t_{i-1}, t_i]$ and from the dynamic programming principle we have,

$$W(t_i, S_{t_i}) = \sup_{(\sigma_t, t\in[t_i,t_{i+1}])} \mathbb{E}\left[W(t_{i+1}, S_{t_{i+1}})\middle|S_{t_i}\right].$$

Then proceeding backwards we can find the optimal control σ_t over each of the intervals $[t_i, t_{i+1}]$ separately. In order to make each of the n optimisation problems tractable we substitute σ_t on the interval $[t_i, t_{i+1}]$ for some parametric function λ_i.
In detail, we proceed as follows ([8], Section 3.2):

1. Simulate N_1 paths $(S^{(p)})_{p=1}^{N_1}$ with some diffusion, say, for instance, the log-normal diffusion with volatility $\hat{\sigma} = \overline{\sigma}$

2. For $i = n-1, n-2, ..., 1, 0$, find a numerical solution θ_i^* of the maximisation problem:

$$\sup_{\theta_i\in\Theta} h(\theta_i), \quad h(\theta_i) = \frac{1}{N_1}\sum_{p=1}^{N_1} V((S_t^{(p)})_{0\le t\le T}), \quad dS_t = \sigma_t S_t dB_t$$

   where

$$\sigma_t = \begin{cases} \hat{\sigma} & \text{if } t \in [0, t_i) \\ \lambda_i(S_t; \theta_i) & \text{if } t \in [t_i, t_{i+1}) \\ \lambda_j(S_t; \theta_j^*) & \text{if } t \in [t_j, t_{j+1}], j \ge i+1 \end{cases}$$

3. Independently, simulate N_2 replicas of S using $\sigma_t = \lambda_i(S_t; \theta_i^*)$ for $t \in [t_i, t_{i+1}]$ and compute $\frac{1}{N_2}\sum_{p=1}^{N_2} V((S_t^{(p)})_{0\le t\le T})$

In practice it can be difficult to find an appropriate parameterisation which allows for speedy and accurate optimisation. However there are a few heuristics that can be applied to make this procedure more robust. To this end we consider first a simple call spread, and then see how this procedure generalises to a more complex cliquet.

Remark 4.3.1. Since we know the optimal control takes values either $\underline{\sigma}$ or $\overline{\sigma}$ we may want to consider parameterisations only of the form,

$$\lambda_i(S_t; \theta_i) = \underline{\sigma} + (\overline{\sigma} - \underline{\sigma})\mathbb{1}_{\gamma_i(S_t;\theta_i)\ge 0}$$

Here γ_i is a parametric function that separates the regions where the gamma at time t is positive or negative.

Remark 4.3.2. The final optimisation is necessary, when we perform the maximisation we are inherently over-fitting the parameters to the initially generated paths, thus to counteract this we re-simulate the paths with the 'optimal' parameters to obtain a less biased estimate of the optimal price.

4.3.2 Call Spread

The remainder of this chapter is my own work. First recall that the payoff of a call spread with strikes $K_d < K_u$, is given by,

$$V(S_T) = (S_T - K_d)^+ - (S_T - K_u)^+$$

Furthermore, notice that in the Black-Scholes setting and analogously in the uncertain volatility setting, there exists a single zero gamma boundary for the call spread payoff. We also know that from the Black-Scholes formula and Chapter 3 that the unique zero gamma boundaries converge to the geometric average of the strikes $\sqrt{K_d K_u}$.
Given that there is a unique zero gamma boundary, instead of solving n optimisations as in the general method in Section 4.3.1, we can just find a parameterisation for the zero gamma boundary. More precisely, we proceed as follows

1. Simulate N_1 paths $(S^{(p)})_{p=1}^{N_1}$ with some diffusion, say, for instance, the log-normal diffusion with volatility $\hat{\sigma} = \overline{\sigma}$
2. Find a numerical solution θ^* of the maximisation problem:

$$\sup_{\theta \in \Theta} h(\theta), \quad h(\theta) = \frac{1}{N_1} \sum_{p=1}^{N_1} V((S_t^{(p)})_{0 \le t \le T}), \quad dS_t = \sigma_t S_t dB_t$$

   where

$$\sigma_t = \underline{\sigma} + (\overline{\sigma} - \underline{\sigma}) \mathbb{1}_{\gamma(t;\theta) < S_t}$$

3. Independently, simulate N_2 replicas of S using $\sigma_t = \underline{\sigma} + (\overline{\sigma} - \underline{\sigma}) \mathbb{1}_{\gamma(t;\theta^*) < S_t}$ and compute $\frac{1}{N_2} \sum_{p=1}^{N_2} V((S_t^{(p)})_{0 \le t \le T})$

For experiments we trial three different parameterisations:

- A constant boundary at the value θ

$$\gamma(t;\theta) = \theta$$

- A linear boundary where θ is the value of the boundary at time 0

$$\gamma(t;\theta) = \theta + \frac{\sqrt{K_d K_u} - \theta}{T} t$$

- A quadratic boundary where θ_1 and θ_2 are the values of the boundary at time 0 and $\frac{T}{2}$ respectively

$$\gamma(t;\theta_1,\theta_2) = \theta_1 + \frac{4\theta_2 - \sqrt{K_d K_u} - 3\theta_1}{T} t + \frac{2\sqrt{K_d K_u} + 2\theta_1 - 4\theta_2}{T^2} t^2$$

Remark 4.3.3. In the case of the linear and quadratic boundaries where we would naturally parameterise with 2 and 3 parameters respectively may lead to a more optimal fit, however for the optimisation it is much quicker and simpler to optimise over a lower dimensional space. Furthermore, the additional information for an end point of these boundaries can improve robustness of the optimisation with respect to initial parameters. Finally the choice to parameterise by points that the curve passes gives the additional benefit of interpretability.

4.3.3 Results

We now draw our attention to Figure 4.3 where we compare the results produced for various call spreads, using the three suggested parametric functions; constant boundary, linear boundary and quadratic boundary. For more information on the Model Comparison quantity, refer to Section 4.1. Generally, all three parametric functions perform well and produce reasonable estimates for the maximal arbitrage-free price for the contract variations. The similarity in results suggests

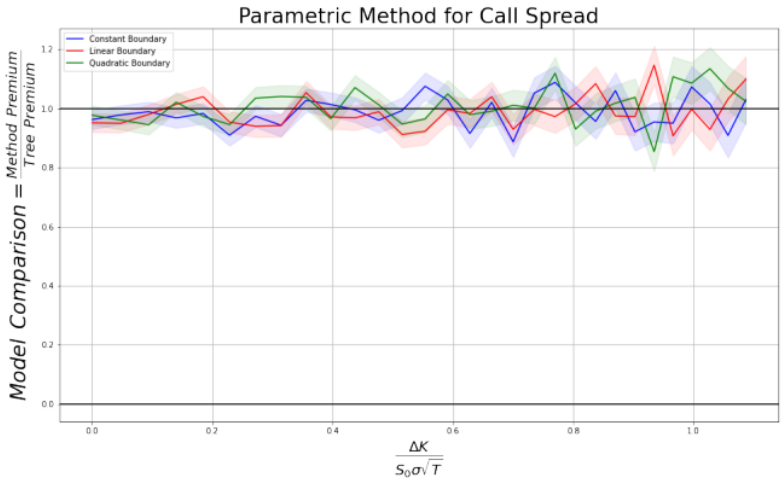

Figure 4.3: Comparison of the parametric method for pricing a variety of call spreads when using different parametric functions for the zero gamma boundary; constant, linear and quadratic. Here we consider a model uncertainty of $\underline{\sigma} = 0.18, \overline{\sigma} = 0.22$, then we take $\sigma = 0.2$ — the mid point

that there is no immediate advantage when using a quadratic parameterisation which requires an additional parameter, and so takes longer to maximise.

We refer to Figure 4.4 for an example of some calibrated boundaries for a call spread with strikes $K_d = 90, K_u = 100$ and $S_0 = 95$. In black we observe the zero gamma boundaries derived from the Black-Scholes formula for a call spread, in blue the constant parameterisation, red the linear parameterisation and green the quadratic parameterisation. Notice, that the calibrated constant boundary in general will not converge to the end point $\sqrt{K_d K_u}$. This raises an important consideration mentioned in Remark 4.3.3, although we know that the end point of the perfectly calibrated boundary should be $\sqrt{K_d K_u}$, since the parameterisation does not adequately capture the shape of the perfectly calibrated boundary, we cannot expect the optimal parameters to yield a curve with the endpoint as $\sqrt{K_d K_u}$.

This raises the question of whether it is wise to fix the endpoint when we consider the linear and quadratic parameterisations. For this we refer to Figure 4.5, where we observe a particular instance of the objective function that we wish to maximise. Recall that $h(\theta)$ corresponds to the estimated price of the derivative when the parametric function defining the boundary has parameter θ. In this example we consider a linear parameterisation, and a range of parameters $\theta \in [90, 100]$. Recall that this parameter specifies the position of the boundary at time $t = 0$, for Figure 4.4 this parameter has value ≈ 93.8. We notice that for this particular example, the objective function for values $\theta \in [90, 97]$ produce a reasonably accurate estimate of the maximal arbitrage-free price (estimated to be ≈ 5.01 from the Trinomial Tree). This has the implication that the curves presented in Figure 4.4 will in general vary upon each run. Thus although it is advantageous to look at methods that model the zero gamma boundary, it is still unclear as to what these boundaries look like in the case of uncertain volatility. From a practical standpoint, this is not a problem, since we still achieve accurate prices.

The method here is specific to a call spread, alone this is not very useful since it would be simpler to use the trinomial tree method. However, developments here motivate a method for pricing cliquets as follows.

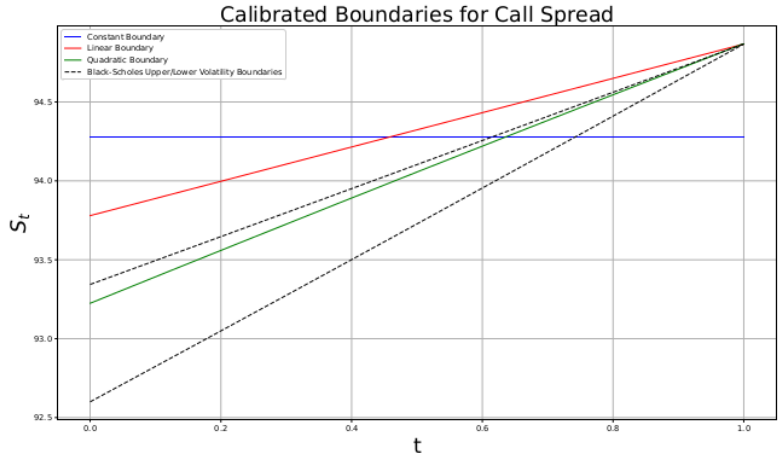

Figure 4.4: Example of calibrated zero gamma boundaries for a call spread with a variety of parametric functions. Here we use $S_0 = 95, K_d = 90, K_u = 100, T = 1 \underline{\sigma} = 0.18, \underline{\sigma} = 0.22$. The black dashed lines correspond to the Black-Scholes zero gamma boundaries with volatilities $\underline{\sigma}$ and $\overline{\sigma}$; hence two curves

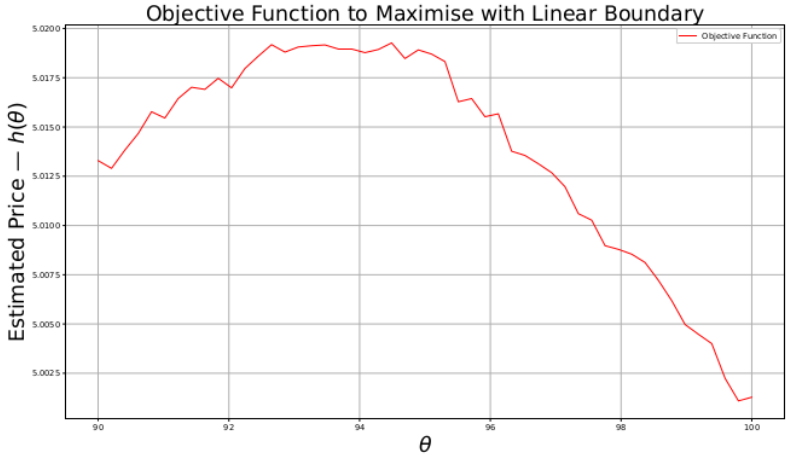

Figure 4.5: Example of objective function (estimate price of call spread) to maximise over the parameter θ in the case of a linear parameterisation for the boundary. Here we use $S_0 = 95, K_d = 90, K_u = 100, T = 1, \underline{\sigma} = 0.18, \underline{\sigma} = 0.22$

Chapter 5

Cliquets

The work presented in this chapter is my own. We begin by introducing cliquets, how their payoff is structured, we use log returns as it makes the mathematics simpler however has minimal effects on the final method. Suppose we have reset periods $T_0, T_1, ..., T_n$, a cap and floor $C > F$ and some strike K, consider the following payoff

$$V((S_t)_{0\leq t\leq T}) = \left(\sum_{i=1}^{n}\left(\log\frac{S_{T_i}}{S_{T_{i-1}}}\right)_F^C - K\right)^+$$

where for shorthand we write for $C > F$

$$(x)_F^C = \max(\min(x, C), F)$$

Here we chose our global feature to be a call with strike K, however the method is also generalisable to a global cap, global floor or even both. Instead of dealing with S_t directly, we consider the log return $X_t := \log S_t$, by Itô's formula we have in the uncertain volatility case

$$dX_t = -\frac{1}{2}\sigma_t^2 dt + \sigma_t dB_t, \quad X_0 = \log S_0 \in \mathbb{R}$$

where $\sigma_t \in [\underline{\sigma}, \overline{\sigma}]$ is non-anticipative. We can define the path dependent variable A_t which is updated discretely at deterministic times,

$$A_t = \begin{cases} (X_t - X_0)_F^C & \text{if } t \leq T_1 \\ \\ A_{T_{i-1}} + (X_t - X_{T_{i-1}})_F^C & \text{if } T_{i-1} \leq t \leq T_i \end{cases} \tag{5.0.1}$$

Then, alternatively the payoff can be written as,

$$V(A_T) = (A_T - K)^+$$

Now notice that the fair value of the derivative at time $t \in [T_{i-1}, T_i)$ depends only on $A_{T_{i-1}}$ and $X_t - X_{T_{i-1}}$ (which for ease of notation we will write the dependence as simply X_t). Furthermore, from the dynamic programming principle, for $t \in [T_{i-1}, T_i)$

$$W(t, X_t, A_{T_{i-1}}) = \sup_{\sigma_r\in[\underline{\sigma},\overline{\sigma}];r\in[t,T_i]} \mathbb{E}\left[W(T_i, X_{T_i}, A_{T_{i-1}})|(X_t, A_{T_{i-1}})\right]$$

Thus as in the derivation of the BSB equation, we can can construct an analogous series of non-linear second order partial differential equations which are solved backwards i.e. on the intervals $[T_{n-1}, T_n], [T_{n-2}, T_{n-1}], ...$ So for $t \in [T_{i-1}, T_i]$

$$\frac{\partial W}{\partial t}(t, x, a) + \frac{1}{2}\sup_{\sigma\in[\underline{\sigma},\overline{\sigma}]}\left[\sigma^2\left(\frac{\partial^2 W}{\partial x^2}(t, x, a) - \frac{\partial W}{\partial x}(t, x, a)\right)\right] = 0 \tag{5.0.2}$$

Where the terminal conditions are given by

$$\begin{cases} W(T_n, x, a) = V(a + (x)_F^C) & \text{for the case } t \in [T_{n-1}, T_n] \\ \\ W(T_i^-, x, a) = W(T_i^+, x, a) & \text{for the case } t \in [T_{i-1}, T_i] \end{cases} \tag{5.0.3}$$

Notice that the optimal control σ_t is again of bang-bang type as expected since there is only one underlying, it is given by

$$\sigma_t = \begin{cases} \overline{\sigma} & \text{if } \frac{\partial^2 W}{\partial x^2} - \frac{\partial W}{\partial x} \geq 0 \\ \\ \underline{\sigma} & \text{if } \frac{\partial^2 W}{\partial x^2} - \frac{\partial W}{\partial x} < 0 \end{cases} \tag{5.0.4}$$

5.1 Extended Trinomial Tree

We now present how ideas in Section 4.2 can be adapted to price the cliquet. We begin with a definition to make the explanation simpler.

Definition 5.1.1 (Trinomial Tree)**.** A tree $\mathcal{T}$ is a collection of nodes and transition probabilities approximating a discrete time version of S_t (the underlying asset) as outlined in Section 4.2. By providing a simple payoff $V(S_T)$, the tree $\mathcal{T}$ can be *solved* by following the algorithm in Section 4.2.1.

We then extend this notion of a tree $\mathcal{T}$ to include additional information necessary for contingent claims on the log return of the underlying asset.

Definition 5.1.2 (Trinomial Tree with Returns)**.** Let $\mathcal{T}$ be a Trinomial Tree, then for each node (i, j) of $\mathcal{T}$ we define the state $X_i^j := \log(S_i^j/S_0)$. We denote this tree with additional information by $\mathcal{T}^{\mathcal{R}}$. The tree $\mathcal{T}^{\mathcal{R}}$ can be *solved* for payoffs of the form $V(S_T, X_T)$ exactly as in Section 4.2.1.

Now the algorithm proceeds as follows, recall we have n reset periods $[T_{i-1}, T_i]$

1. We create $1+N+N^2+...+N^{n-1}$ Trinomial Trees with Returns. We group them in n levels, so for level $i \in \{0, 1, ..., n-1\}$ trees are indexed via the multi-index $(j_1, ..., j_i) \in \{1, ..., N\}^i$. The tree is denoted by $\mathcal{T}^{\mathcal{R}}(j_1, ..., j_i)$. The case where $i = 0$, the tree is denoted by $\mathcal{T}^{\mathcal{R}}(0)$. The idea here is that from each tree at level i, N more trees stem from each of the final nodes

2. For each of the trees $\mathcal{T}^{\mathcal{R}}(j_1, ..., j_i)$, denote the terminal log returns by $X_k(j_1, ..., j_i) := X_N^k$ for $k \in \{1, ..., N\}$. Similarly $X_k(0)$ corresponds to the terminal log returns of the tree $\mathcal{T}^{\mathcal{R}}(0)$. Notice these values are independent of what we choose as our initial S_0 for each of the trees. We introduce $1 + N + N^2 + ... + N^{n-1}$ auxiliary variables $A(j_1, ..., j_i)$ for each of the trees defined as follows

$$\begin{aligned} A(0) &= 0 \\ A(j_1) &= (X_{j_1}(0))_F^C \quad \text{for } j_1 \in \{1, ..., N\} \\ A(j_1, ..., j_i) &= A(j_1, ..., j_{i-1}) + X_{j_i}(j_1, ..., j_{i-1}) \end{aligned}$$

   These auxiliary variables are simply the discrete case of the variables defined in (5.0.1), however only at the times $T_0, T_1, ..., T_{n-2}$

3. Now the idea is to solve each tree starting from level $n-1$, then $n-2$ and so on until we have solved $\mathcal{T}_0^{\mathcal{R}}$ by using the terminal conditions analogous to (5.0.3). More precisely, for the trees $\mathcal{T}^{\mathcal{R}}(j_1, ..., j_{n-1})$ we solve with the payoff defined by

$$V(X_T) = \left((X_T)_F^C + A(j_1, ..., j_{n-1}) - K\right)^+$$

   We denote the solution of tree $\mathcal{T}^{\mathcal{R}}(j_1, ..., j_i)$ by $W(j_1, ..., j_i)$. Then for trees $\mathcal{T}^{\mathcal{R}}(j_1, ..., j_i)$ we solve with the payoff defined by

$$V(X_T) = \sum_{k=1}^{N} W(j_1, ..., j_i, k) \mathbb{1}_{X_T = X_k(j_1, ..., j_i)}$$

   We keep going until we find $W(0)$ which is the solution.

Remark 5.1.3. The scheme presented here is simply an explicit scheme for solving the series of PDEs defined in (5.0.2)

Remark 5.1.4. The number of trees we have to solve with the number of periods defining the cliquet increases exponentially, thus the scheme is only feasible for up to 3 reset periods, regardless it is useful when benchmarking the parametric methods for simple cases

Remark 5.1.5. There are many optimisations that can be made to the algorithm, for example, for a given level i there are multiple unnecessary repeats of the auxiliary variable leading to the same sequence of trees being solved. By grouping these repeats and only solving unique variations of the trees one can imagine much of the computational complexity would be reduced. However the exact groupings depend on the underlying parameters for $\overline{\sigma}, C, F$ and K

5.1.1 Evaluation

Since this method will be our benchmark, it is important to understand how large N must be to achieve convergence. In practice it is only feasible to apply the trinomial tree for up to 3 periods, additionally, the computational speed begins to fall off as $N > 100$. We refer to Figure 5.1 for a similar analysis as was performed in Section 4.2.2. By setting $\underline{\sigma} = \overline{\sigma} = \sigma$, the Trinomial Tree method should return to us the Black-Scholes price of a cliquet. Furthermore, it is straight forward to price the cliquet in Monte Carlo, simply note that we have the following equality in distribution

$$X_{T_i} - X_{T_{i-1}} = -\frac{1}{2}\sigma^2 (T_i - T_{i-1}) + \sigma\sqrt{T_i - T_{i-1}}Z$$

where Z follows a standard normal distribution. With this quantity we can simulate the returns over each period and compute the price in a Monte Carlo fashion. Then Figure 5.1 illustrates the percentage difference between the price produced from the trinomial tree and the Monte Carlo. Here we produce prices where $F = -0.01, C = 0.04, \sigma = 0.2$ and $T = 1$ with 3 reset periods. The dashed lines indicate the boundary for an error of $\pm 1\%$, of which we observe that we require $N = 100$ to remain somewhat consistently within this boundary. Nevertheless, as the strike K becomes too large the trinomial tree starts to form inaccuracies.

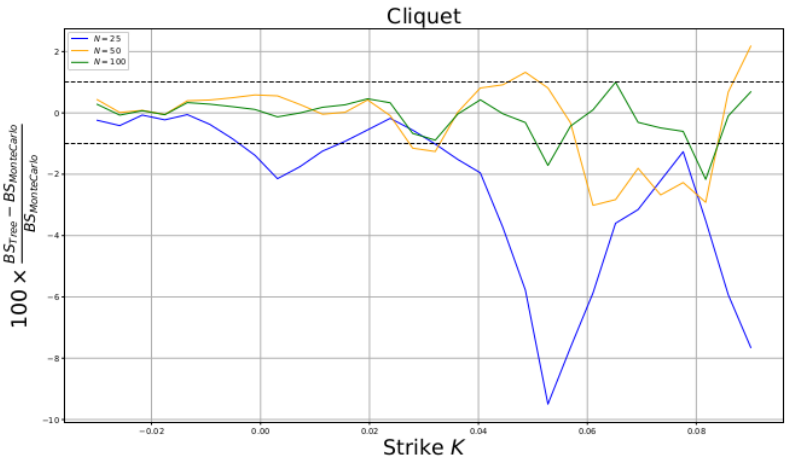

Figure 5.1: Numerical error of the trinomial tree (Black-Scholes) for cliquets. The y-axis is the percentage difference between the trinomial tree method and Monte Carlo method. The parameters used for the 3-period cliquet are $F = -1\%, C = 4\%, T = 1, \sigma = 20\%$

5.2 Parametric Method for Cliquets

5.2.1 Methodology

Pricing when $t \in [T_{n-1}, T_n]$

We begin by studying the price on the final period $[T_{n-1}, T_n]$. Conditioning on $A_{T_{n-1}}$ and defining the return given by the asset over period $[T_{n-1}, T_n]$ by $R_n = X_{T_n} - X_{T_{n-1}}$, the payoff is given by

$$\left((R_n)_F^C - \left(K - A_{T_{n-1}}\right)\right)^+ = \begin{cases} (R_n)_F^C - (K - A_{T_{n-1}}) & \text{if } F \geq K - A_{T_{n-1}} \\ \left(R_n - (K - A_{T_{n-1}})\right)^+ - (R_n - C)^+ & \text{if } C \geq K - A_{T_{n-1}} \geq F \\ 0 & \text{if } K - A_{T_{n-1}} \geq C \end{cases}$$

In the first case, we simply have a call-spread payoff on R_n with a translation upwards/downwards, the second is also a call-spread with strikes $K_d = K - a$ and $K_u = C$. The final payoff is 0 since it is impossible to obtain a final payoff when the value of $A_{T_{n-1}}$ is too small, note that in this case it does not matter what our optimal control for σ_t is, since either way the payoff is 0.

A visual example of this analysis is provided in Figure 5.2. We can see from the figure that in the case where the strike is 0%, we are always in case 1 and 2. Note that the thickness of each group is a consequence of conditioning on $A_{T_{n-1}}$ over a small range of values, rather than a single value. We could have chosen smaller buckets, however this would require more paths in the simulation to observe the structure of the payoffs clearly. So we confirmed over the final period, conditioned on $A_{T_{n-1}}$, we observe a call-spread 'like' payoff. Thus we can find the maximal arbitrage-free price on this period in a similar fashion to Section 4.3.2.

The x-axis position of the vertices can be computed as a function of $A_{T_{n-1}}$, and is given by

$$\begin{cases} F \text{ and } C & \text{if } F \geq K - A_{T_{n-1}} \\ (K - A_{T_{n-1}}) \text{ and } C & \text{if } C \geq K - A_{T_{n-1}} \geq F \end{cases} \tag{5.2.1}$$

in the case $K - A_{T_{n-1}} \geq C$ the payoff is 0, so there are no vertices. The idea here is to find, similarly to the call-spread, a 'zero gamma boundary' over this final period. In fact this is no longer a path, since this boundary will now be a function of t and $A_{T_{n-1}}$. More specifically, it will be a function $\gamma : [T_{n-1}, T] \times [(n-1)F, (n-1)C] \to \mathbb{R}$. Then over the period $[T_{n-1}, T_n]$, we will simulate the variable X_t according to the stochastic differential equation

$$dX_t = -\frac{1}{2}\left(\underline{\sigma} - (\overline{\sigma} - \underline{\sigma})\mathbb{1}_{X_t \leq \gamma(t, A_{T_{n-1}})}\right)^2 dt + \left(\underline{\sigma} - (\overline{\sigma} - \underline{\sigma})\mathbb{1}_{X_t \leq \gamma(t, A_{T_{n-1}})}\right) dB_t, \quad X_{T_{n-1}} = 0$$

i.e. $\underline{\sigma}$ if we are above the boundary and $\overline{\sigma}$ if we are below. Then by setting $R_n = X_{T_n}$, we compute the payoff and thus the price by averaging over multiple paths. The next problem is the choice of the parametric function $\gamma(t, A_{T_{n-1}})$.

For a given $A_{T_{n-1}}$, we know the payoff is identical to a call spread, our analysis in Section 4.3.2 suggests that a simple affine function, or even a constant function $t \mapsto \gamma(t, A_{T_{n-1}})$ will suffice to produce accurate results. In fact, we might be tempted to also conclude that the boundary should converge to the geometric average between the two 'strikes' — the two vertices, which in general vary with $A_{T_{n-1}}$. However, recall that the boundary defining the optimal control (5.0.4), is not just dependent on the second order derivative, but also on the first order derivative i.e.

$$\frac{\partial^2 W}{\partial x^2} - \frac{\partial W}{\partial x} = 0$$

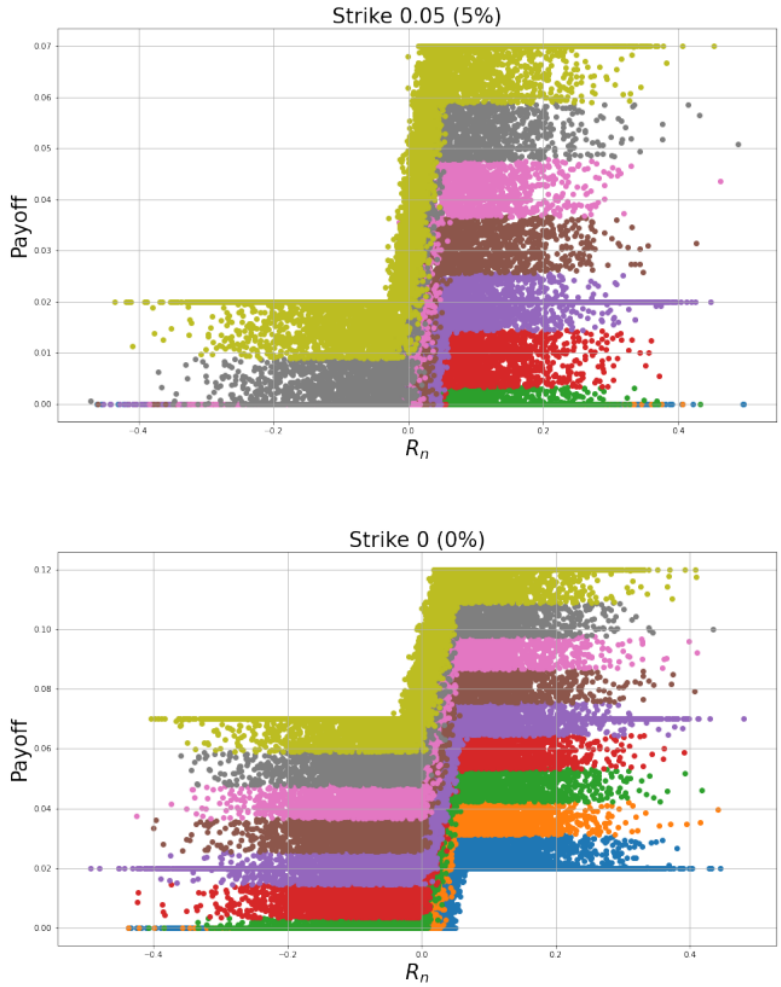

Figure 5.2: Simulated payoffs for the cliquet as a function of returns R_n over period $[T_{n-1}, T_n]$, conditioned on $A_{T_{n-1}}$ (grouped in buckets represented by the colours, $a_i \leq A_{T_{n-1}} \leq a_{i+1}$ for some small $a_{i+1} - a_i$). Values used: $C = 4\%$, $F = -1\%$, $T = 1$, $n = 4$ (3 periods), strikes 5% and 0%

Notice that the payoff, conditioned on $A_{T_{n-1}}$ is always increasing in R_n, so the price over this period, at time t, will also be increasing in X_t, thus we have $\frac{\partial W}{\partial x}(t, X_t) > 0$. The consequence here, is that the boundary must always lie in some strictly convex region of the price as a function of X_t. We must also justify that there is still a single boundary where the above equality holds, however this is fairly intuitive taking into account the shape of the call-spread payoff

- Notice that $x \to \frac{\partial^2 W}{\partial x^2}(t, x)$ is a **strictly** decreasing function, furthermore $x \to \frac{\partial W}{\partial x}(t, x)$ is increasing when $\frac{\partial^2 W}{\partial x^2}(t, x) \geq 0$, and decreasing when $\frac{\partial^2 W}{\partial x^2}(t, x) \leq 0$.

- So we have that $\frac{\partial W}{\partial x}(t, x)$ obtains a maximum at the **unique** point where $\frac{\partial^2 W}{\partial x^2}(t, x) = 0$. Thus for x in the region such that $\frac{\partial^2 W}{\partial x^2}(t, x) \geq 0$, since $\frac{\partial W}{\partial x}(t, x)$ is increasing, the function

$$x \to \frac{\partial^2 W}{\partial x^2}(t,x) - \frac{\partial W}{\partial x}(t,x) \quad \text{is decreasing}$$

- Next, notice that for all x, we have $\frac{\partial W}{\partial x}(t,x) \geq 0$. Finally, for all x where $\frac{\partial^2 W}{\partial x^2}(t,x) < 0$

$$\frac{\partial^2 W}{\partial x^2}(t,x) - \frac{\partial W}{\partial x}(t,x) < -\frac{\partial W}{\partial x}(t,x) \leq 0$$

- Thus we can conclude that for a given t, there is a single point x such that

$$\frac{\partial^2 W}{\partial x^2}(t,x) - \frac{\partial W}{\partial x}(t,x) = 0$$

  Note that this analysis also shows that this unique point lies in a region where $\frac{\partial^2 W}{\partial x^2}(t,x) > 0$ as previously mentioned

Recall that the process X_t defining R_n for the returns over $[T_{n-1}, T_n]$ satisfy,

$$dX_t = -\frac{1}{2}\sigma_t^2 dt + \sigma_t dB_t$$

We see that there is a strictly negative drift, implying that the final distribution R_n satisfies

$$\mathbb{E}\left(R_n \middle| X_{T_{n-1}}\right) = \mathbb{E}\left(X_{T_n} - X_{T_{n-1}} \middle| X_{T_{n-1}}\right) = \mathbb{E}\left(-\frac{1}{2}\int_{T_{n-1}}^{T_n} \sigma_t^2 dt\right) < 0$$

In words we expect R_n to be negative, of course this is an artefact of our choice that the risk-free rate is 0, nevertheless this analysis provides intuition for what the optimal boundary should be. It follows that the optimal strategy is less 'risky' than in Section 4.3.2 where the underlying asset in the call spread was a martingale.

Although we might have a fair understanding of how the parametric function $\gamma(\cdot)$ depends on t, we still need to understand how it depends on $A_{T_{n-1}}$ if we are to make a suitable choice. Recall from Section 4.3.2 that the optimal boundary was close to the geometric average of the two strike throughout the whole time horizon. If we assume the same for each of the curves conditioned on $A_{T_{n-1}}$ in Figure 5.2, it may be useful to understand how the middle point of the vertices in (5.2.1) change as we vary $A_{T_{n-1}}$. Since the vertices can be negative numbers, and the geometric average is not necessarily relevant as we are no longer looking at the gamma of the option, we simply focus on the arithmetic average instead.

In Figure 5.3 we see the general dependence of the arithmetic average of the vertices (5.2.1), colour coded based on which case we are in. In the case $A_{T_{n-1}} \geq K - F$, we will always be in-the-money of the global call feature, thus we only have a translation (up and down) of the call-spread 'like' payoff, thus the middle point of the vertices should not change. Intuitively, neither should our strategy; one could argue that the 'game' has not changed, just the *scale* of the rewards. The second case, $C \geq K - A_{T_{n-1}} \geq F$, we note that the right-most vertex stays fixed at C, whilst the left-most converges to C as $A_{T_{n-1}}$ decreases. The point at which we should switch the volatility should scale as one might imagine with this; consider the simple case with the standard call spread, as the strike K_d converges to K_u, the geometric average also converges to K_u. Finally, we do not consider the case where $A_{T_{n-1}} \leq K - C$, since here we are always out-of-the-money no matter what happens.

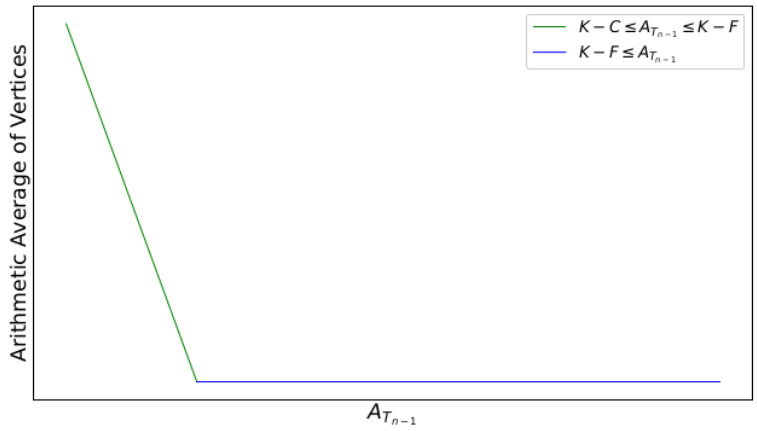

Figure 5.3: Example of the dependence of the parametric boundary on the variable $A_{T_{n-1}}$ over the final period $[T_{n-1}, T_n]$

Now that we have a more complete understanding of what the payoff looks like, and more importantly what the boundary defining the optimal strategy (5.0.4) should look like, we suggest a few different parameterisations for this boundary over the period $[T_{n-1}, T_n]$.

1. A time-homogeneous, piece-wise affine boundary

$$\gamma\left(t, A_{T_{n-1}}; \theta_1, \theta_2\right) = \begin{cases} \theta_1 & \text{if } K - F \leq A_{T_{n-1}} \\ \frac{(K-F)\theta_2 - (K-C)\theta_1 + (\theta_1 - \theta_2)A_{T_{n-1}}}{C-F} & \text{if } K - C \leq A_{T_{n-1}} \leq K - F \\ 0 & \text{o/w} \end{cases}$$

   The parameterisations are chosen so that θ_1 is the value over the period $K - F \leq A_{T_{n-1}}$ and θ_2 is the value of the function at $A_{T_{n-1}} = K - C$. Here we are directly modelling the curve given in Figure 5.3, however note that the purpose of the figure was to provide an idea of the dependence, the actual dependence on $A_{T_{n-1}}$ may not look like this. We also assume here that the boundary does not change over time, for smaller reset periods this is a reasonable approximation.

2. A time-homogeneous, one-parameter fractional boundary

$$\gamma\left(t, A_{T_{n-1}}; \theta_1\right) = \begin{cases} F + \theta_1 (C - F) & \text{if } K - F \leq A_{T_{n-1}} \\ \left(K - A_{T_{n-1}}\right) + \theta_1 \left(C - \left(K - A_{T_{n-1}}\right)\right) & \text{if } K - C \leq A_{T_{n-1}} \leq K - F \\ 0 & \text{o/w} \end{cases}$$

   The idea here is that θ_1 represents where the boundary lies within the two vertices of the payoff i.e. if $\theta_1 = 0.5$, then the boundary lies half way between the payoffs

The parameterisations we consider here are of course not exhaustive, and in no means will they perfectly describe the boundaries. However the idea here is to model the boundary as best as we can with as few parameters as possible, and it turns out these parameterisations do a good job of this. There are many variations we could consider of the above parameterisations by introducing more parameters, especially in the fractional boundary case.

The discussion so far has been centred around pricing in the final period $[T_{n-1}, T_n]$, which is not quite so impressive since we could actually just use the tree approach to price here effectively. The idea is to now generalise this approach to finding this maximal arbitrage-free price at time $t = 0$. In short this will involve parameterising a separate boundary for each of the reset periods, however it will turn out that the boundary in these previous period are still fairly intuitive.

Pricing when $t = 0$

The problem here is that the payoff at time $t \in [T_{i-1}, T_i]$ not only depends on the return within this period, but also on returns accumulated over the following periods. The idea is that nevertheless, the price at time $t \in [T_{i-1}, T_i]$ conditioned on $A_{T_{i-1}}$ as a function of $X_t - X_{T_{i-1}}$ will still give a similar shape with strictly one convex region and one concave region, thus a single boundary defining where we should choose $\underline{\sigma}$ or $\overline{\sigma}$. The rationale here is that the returns over the future periods, do not depend on the return over the current period. More generally, setting $R_j = X_{T_j} - X_{T_{j-1}}$ the payoff is given by

$$\left(\left[(R_i)_C^F \right] + \left[\sum_{j=i+1}^{n} (R_j)_C^F \right] - \left[K - A_{T_{i-1}} \right] \right)^+$$

Focusing on pricing for $t \in [T_{i-1}, T_i]$, the payoff can be grouped into the above three terms. Notice the similarity here with the payoff in the final period, in-fact it is still a call-spread 'like' payoff, however there is still a random term, the sum of the future returns

$$B_{T_i} := \left[\sum_{j=i+1}^{n} (R_j)_C^F \right]$$

Notice that each R_j are i.i.d, importantly this sum, B_{T_i}, is independent of R_i. So we can approach this in an identical fashion as the final period; the payoff can be expressed as

$$\left((R_i)_F^C - \left(K - A_{T_{i-1}} - B_{T_i} \right) \right)^+ = \begin{cases} (R_n)_F^C - (K - A_{T_{n-1}} - B_{T_i}) & \text{if } F \geq K - A_{T_{n-1}} - B_{T_i} \\ \left(R_n - (K - A_{T_{n-1}} - B_{T_i}) \right)^+ - (R_n - C)^+ & \text{if } C \geq K - A_{T_{n-1}} - B_{T_i} \geq F \\ 0 & \text{if } K - A_{T_{n-1}} - B_{T_i} \geq C \end{cases}$$

So for a given realisation of B_{T_i}, we are again in the case of looking at call-spreads, in fact we can write for the price at time $t \in [T_{i-1}, T_i]$

$$\begin{aligned} W(t, X_t, A_{T_{i-1}}) &= \sup \mathbb{E} \left[\left((R_i)_F^C - \left(K - A_{T_{i-1}} - B_{T_i} \right) \right)^+ \middle| \left(X_t, A_{T_{i-1}} \right) \right] \\ &= \sup \mathbb{E} \left[\mathbb{E} \left[\left((R_i)_F^C - \left(K - A_{T_{i-1}} - B_{T_i} \right) \right)^+ \middle| \left(B_{T_i}, A_{T_{i-1}} \right) \right] \middle| \left(X_t, A_{T_{i-1}} \right) \right] \end{aligned}$$

where the inner expectation is precisely a call-spread 'like' payoff, thus we end up averaging over these payoffs with respect to the distribution of B_{T_i}. The question is then if the distribution of B_{T_i} is 'nice' enough that this average ends up maintaining the shape of the price of a call-spread, more precisely, if we end up maintaining a single boundary separating the regions of volatility $\underline{\sigma}$ and $\overline{\sigma}$.

For a given time t and value $A_{T_{i-1}}$, we denote the unique boundary as a function $B_{T_i} \mapsto z\left(B_{T_i}\right)$, then the boundary as a function of $(t, A_{T_{i-1}})$ is given by averaging over the distribution of B_{T_i}. However, since the values that B_{T_i} takes, lie in the compact set $[(n-i)F, (n-i)C]$, this averaging gives us a unique point that lies in the compact set

$$\left\{ z\left(B_{T_i}\right) \mid B_{T_i} \in [(n-i)F, (n-i)C] \right\}$$

Additionally, since $(t, A_{T_{i-1}})$ is independent of B_{T_i} we obtain a unique continuous boundary as in the case of the final period. We now need only model this boundary as we did in the in the case

of the final period. In fact, since the boundary is an averaging independent of $(t, A_{T_{i-1}})$, and the payoff structure in the period $[T_{i-1}, T_i]$ is similar to the final period, we expect a similar boundary, so we may trial the same parameterisations.

5.2.2 The Algorithm

We now bring all of the ideas above together to provide an algorithm for finding the maximal arbitrage-free price of a Cliquet where the volatility is uncertain. So, for some parametric functions $(\gamma_i)_{i=1}^n$

1. Simulate N_1 paths $\left(X^{(p)}\right)_{p=1}^{N_1}$ with the diffusion with volatility $\hat{\sigma} = \overline{\sigma}$

$$dX_t = -\frac{1}{2}\hat{\sigma}^2 dt + \hat{\sigma} dB_t, \quad X_0 = 0$$

   Set $R_i^{(p)} = X_{T_i}^{(p)} - X_{T_{i-1}}^{(p)}$ for the returns over each period, furthermore set $A_i^{(p)} = \sum_{j=1}^{i} \left(R_j^{(p)}\right)_F^C$ for the rolling capped/floored returns. The payoff is then given as a function of $A_n^{(p)}$.

2. For $i = n, n-1, ..., 1$, find a numerical solution θ_i^* of the maximisation problem:

$$\sup_{\theta_i \in \Theta} h(\theta_i), \quad h(\theta_i) = \frac{1}{N_1} \sum_{p=1}^{N_1} V(A_n^{(p)}),$$
$$dX_t = -\frac{1}{2}\sigma_t^2 dt + \sigma_t dB_t, \quad X_0 = 0$$

   where

$$\sigma_t = \begin{cases} \hat{\sigma} & \text{if } t \in [0, T_{i-1}) \\ \underline{\sigma} + (\overline{\sigma} - \underline{\sigma}) \mathbb{1}_{X_t - X_{T_{i-1}} \leq \gamma_i\left(t, A_{T_{i-1}}; \theta_i\right)} & \text{if } t \in [T_{i-1}, T_i) \\ \underline{\sigma} + (\overline{\sigma} - \underline{\sigma}) \mathbb{1}_{X_t - X_{T_{j-1}} \leq \gamma_j\left(t, A_{T_{j-1}}; \theta_j^*\right)} & \text{if } t \in [T_{j-1}, T_j), j \geq i+1 \end{cases}$$

3. Independently, simulate N_2 replicas of X using $\sigma_t = \underline{\sigma} + (\overline{\sigma} - \underline{\sigma}) \mathbb{1}_{X_t - X_{T_{i-1}} \leq \gamma_i\left(t, A_{T_{i-1}}; \theta_i^*\right)}$ for $t \in [T_{i-1}, T_i]$ and compute $\frac{1}{N_2} \sum_{p=1}^{N_2} V(A_n^{(p)})$

Remark 5.2.1. The algorithm requires n optimisations for each reset period defining the cliquet, this can increase the time complexity and effectiveness of the algorithm drastically if the parametric functions are not chosen with care. This is precisely why it is preferable to choose parametric functions with few parameters. Using the analysis in our methodology helps us obtain these parametric functions.

Notice in Figure 5.3, the vertex splitting the two affine functions is simply given by $K - F$. More generally for the period $[T_{i-1}, T_i]$, one can imagine the vertex is given by approximately $K - F - \mathbb{E}(B_{T_i})$. So if we are in a situation where $\mathbb{E}(B_{T_i}) \approx 0$ we can consider the modified algorithm where we use just a single parametric function $\gamma(\cdot)$. So replacing Step 2 in the algorithm with

- Find a numerical solution θ^* of the maximisation problem:

$$\sup_{\theta_i \in \Theta} h(\theta_i), \quad h(\theta_i) = \frac{1}{N_1} \sum_{p=1}^{N_1} V(A_n^{(p)}),$$
$$dX_t = -\frac{1}{2}\sigma_t^2 dt + \sigma_t dB_t, \quad X_0 = 0$$

  where

$$\sigma_t = \underline{\sigma} + (\overline{\sigma} - \underline{\sigma}) \mathbb{1}_{X_t - X_{T_{i-1}} \leq \gamma\left(t, A_{T_{i-1}}; \theta^*\right)} \quad \text{for } t \in [T_{i-1}, T_i]$$

Here we are reducing the number of optimisations required to just one, clearly in practice this is preferred. However we should evaluate when this modified algorithm is effective since we are making some large assumptions about the similarity of the boundary over different reset periods. On the other hand, at any given period $[T_{i-1}, T_i]$, to get the largest payoff it is always preferable to obtain the largest return over each period. Thus one would expect that the optimal strategy for σ_t to achieve the best payoff should not differ largely.

5.2.3 Results

We now draw our attention to Figure 5.4 where we observe our results for a variety of cliquets with differing strikes K. In our results we only consider 3 periods, since in this case we can use the trinomial tree method to compare against. In blue and red we use the Fractional and Piece-wise affine parameterisations for period defining the cliquet, thus 3 separate optimisations. In green we have our results for the modified algorithm as discussed at the end of Section 5.2.2. For this single parameterisation we use the same Piece-wise affine boundary. In our results we see that the red and blue plots produce similar results, thus in this case the fractional boundary is preferred as it requires just one parameter, thus the maximisation is quicker. We see that the modified algorithm (green) produces slightly better results; this supports our hypothesis that the optimal strategy for σ_t should remain the same no matter which period you are in.

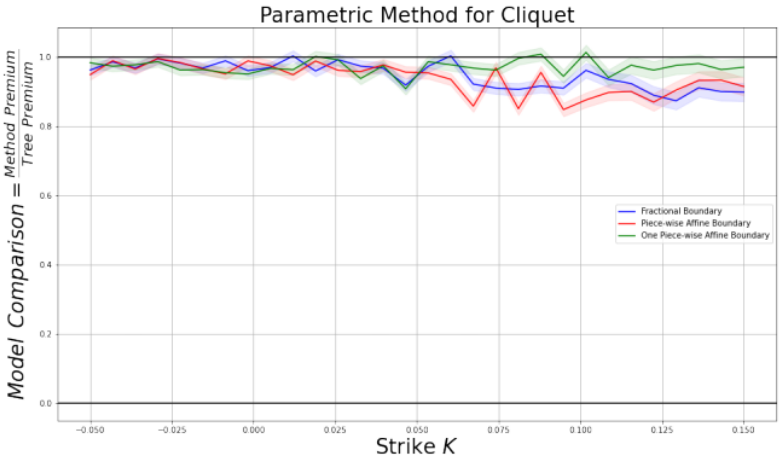

Figure 5.4: Comparison of the parametric method for pricing a variety of cliquets. Parametric functions used are the Fraction parameterisation (blue), the piece-wise affine parameterisation (red), we also look at the modified approach using one piece-wise affine parameterisation for all periods (green). Here we consider a 3-period cliquet with $F = -1\%, C = 4\%, T = 1, \underline{\sigma} = 15\%, \overline{\sigma} = 20\%$ and the strike $K \in [-3\%, 9\%]$

The single optimisation for the entire duration of the options life can also help the maximisation procedure. We previously observed in Section 4.3.3 that the final price of the option can become *insensitive* to the parameters of the parametric function, especially around the true maximum. When performing n optimisations for each period, this effect can only be amplified since tweaking the parameters only has an effect on the returns produced over this period. This problem could grow as the number of reset periods grows, thus yielding another reason in favour of the single parametric function for all periods.

Notice in our results, we only consider strikes K up to 9%, this is primarily because of the inaccuracies we saw from the trinomial tree in Section 5.1.1 as the strike grew too large. From a practical perspective, notice that the true maximal arbitrage-free price will always be an upper bound on

any of the prices produced by these methods (as long as your strategy makes use of an admissible $\sigma_t \in [\underline{\sigma}, \overline{\sigma}]$). So in practice when you do not have a reference price to compare against, you can simply trial multiple different parametric functions and choose the one which gives you the largest price (subject to the variance of this price being relatively small). This is the same approach to model selection used by Guyon and Labordère [8].

In conclusion, the modified algorithm produces successful results, with a time complexity similar to that of the simple call spread presented in Section 4.3.2. Some care should be taken when applying the modified algorithm; if one expects the optimal strategy to change across reset periods, this algorithm is unlikely to produce accurate results. In this setting, it is more suitable to model σ_t on each of the periods individually, as with our original algorithm. However, performing an optimisation on each period can become computationally expensive.

Although we only experimented with a global call feature, the method is extensible to at least global caps/floors, or both. However, when considering alternative global features, one must be careful how to choose the parametric function. We did not experiment with time varying parameterisations for the cliquet; it would be interesting to see the effect of these additional parameters, however it seems this is only feasible at the moment when considering the modified algorithm, else there may be too many parameters to optimise over for each period.

Conclusion

We began our discussions with an introduction to the setting of uncertain volatility. We saw that, here, we can obtain multiple arbitrage-free prices. From this, the maximal arbitrage-free price can be obtained theoretically via the Black-Scholes-Barenblatt PDE. We then moved on to discuss results by Denis and Martini [3] which linked this maximal arbitrage-free price and the cheapest super-replicating price, the property we are actually interested in.

In practice, it is not trivial to deduce the maximal arbitrage-free price. We discussed existing methods presented by Avellaneda et al [1] and Guyon and Labordère [8] which have their limitations. We developed these existing methods by introducing the notion of 'zero gamma boundaries', which led to a robust pricing method for a variety of cliquets. In particular, in our modified algorithm we produced accurate prices with just a single optimisation as oppose to arbitrarily many as suggested in literature.

Despite its success, the developed methods are not as generalisable as those presented in literature. In particular, the method assumes that there is a single boundary defining where we should choose the maximum or minimum volatility. This is not the case for most derivatives, possibly not for more complex global features with cliquets. For further work, it would be interesting to see how one can generalise the method of modelling the boundaries; in particular where there are a fixed number of boundaries, you could parameterise each of these separately and simulate with all of these boundaries together.

The uncertain volatility model considers a family of models where the volatility lies in a certain range. However, one might question how realistic some of the models in the family are. In fact, the optimal strategy for the volatility itself is an unlikely event; namely the volatility switching between two values that may be arbitrarily far apart is an unlikely model. It may be more useful to consider a more specific, yet still unspecified setting where the family of models are more reasonable — perhaps defined under an additional assumption on the continuity (or regularity) of the volatility. With additional restrictions on the volatility, we would surely find that the maximal arbitrage-free price is lower than what we would obtain with the original uncertain volatility model. It would be useful to quantify how much lower this price is with respect to the additional restrictions on the volatility.

Bibliography

[1] M. Avellaneda, A. Levy, and A. Parás. Pricing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance*, 2(2):73–88, 1995.

[2] Freddy Delbaen. Representing martingale measures when asset prices are continuous and bounded. *Mathematical Finance*, 2(2):107–130, 1992.

[3] Laurent Denis and Claude Martini. A theoretical framework for the pricing of contingent claims in the presence of model uncertainty. *The Annals of Applied Probability*, 16(2), may 2006.

[4] Nicole El Karoui and Marie-Claire Quenez. Dynamic programming and pricing of contingent claims in an incomplete market. *SIAM journal on Control and Optimization*, 33(1):29–66, 1995.

[5] Hans Föllmer and Yuri M Kabanov. Optional decomposition and lagrange multipliers. *Finance and Stochastics*, 2(1):69–81, 1997.

[6] Jean-Pierre Fouque and Ning Ning. Uncertain volatility models with stochastic bounds. *SIAM Journal on Financial Mathematics*, 9(4):1175–1207, 2018.

[7] Jim Gatheral. *The volatility surface: a practitioner's guide*. John Wiley & Sons, 2006.

[8] Julien Guyon and Pierre Henry-Labordère. Uncertain volatility model: A monte-carlo approach. 2010.

[9] Dmitrij O Kramkov. Optional decomposition of supermartingales and hedging contingent claims in incomplete security markets. *Probability Theory and Related Fields*, 105(4):459–479, 1996.

[10] Matthieu Leblanc and Claude Martini. Unbounded volatility in the uncertain volatility model. *SSRN Electronic Journal*, 11 2000.

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

[12] Claude Martini. The uncertain volatility model and american options. 01 1999.

[13] Andrea Pascucci. *PDE and martingale methods in option pricing*. Springer Science & Business Media, 2011.