UPPSALA
UNIVERSITET

U.U.D.M. Project Report 2014:14

Hedging of Volatility

Ty Lewis

Examensarbete i matematik, 30 hp
Handledare och examinator: Maciej Klimek
Maj 2014

Department of Mathematics
Uppsala University

Uppsala University

Masters Thesis

Hedging of Volatility

Author:
Ty Lewis

Supervisor:
Maciej Klimek

A thesis submitted in fulfilment of the requirements for the degree of Master of Financial Mathematics

in the

Department of Mathematics

May 21, 2014

"An investment in knowledge pays the best interest."

- Benjamin Franklin

UPPSALA UNIVERSITY

Abstract

Department of Mathematics

Master of Financial Mathematics

Hedging of Volatility

by Ty Lewis

This paper investigates pricing and replication of volatility derivatives; beginning with variance and volatility swaps, moving on to options on those swaps and finally examines a method for pricing VIX options. Volatility derivatives are important tools which allow investors to either hedge the volatility risk in a portfolio or speculate on a deviation of realized volatility from the current market expectation. Since the mid 1990's, volatility derivatives have gained popularity because of their ability to replace a complex portfolio with a single instrument. We investigate the work of Neuberger, [2], Dupire [3], Carr and Madan [4] and many others who contributed to a general understanding of volatility trading and especially variance swap pricing. Later, the methods demonstrated by Carr-Lee [11] and Friz-Gatheral [16] for robust pricing and replication of volatility swaps are examined with simulations. Finally we examine a method that Carr-Lee [13] propose for pricing options on variance and volatility swaps as well as the VIX based on the assumption that realized volatility is lognormally distributed as proposed by Friz-Gatheral [16]. We investigate the performance of this method in a small market study and show that the method produces large discrepancies between theoretical and market prices for out of the money calls and discuss the sources of these discrepancies.

Acknowledgements

Special thanks to Jonas Hägg, Pierre Bäcklund and the ORC team for their insight and support. And to Maciej Klimek, without whose knowledge and support as teacher and supervisor this would not have been possible.

Contents

List of Figures

List of Tables

Chapter 1

Background

1.1 Volatility derivatives

The volatility of an asset's price process is a way of describing its dispersion, or how much it fluctuates. Thus volatility can intuitively be thought of as a metric for describing how "risky" an asset is in the sense that it describes how much the price process is likely to deviate from its expected value. Recognizing this risk, it is fair to assume that investors will demand a premium for holding a volatile asset, motivating the fundamental trade-off between risk and expected return. From the inception of Modern Portfolio Theory resulting from the influential work of Markowitz [1], volatility has played a central role in constructing asset portfolios. But it was not until the 1990's that volatility began to be viewed as an asset class of its own. As a result of the work of Neuberger, [2], Dupire [3], Carr and Madan [4] and many others, the dynamics and pricing of volatility became better understood and trading volumes in volatility derivatives have increased to the present day.

There are many reasons why an investor might want to trade volatility. For example, she may wish to speculate that the volatility of an asset will increase or hedge the volatility risk in a given portfolio. Standard derivatives such as European options provide only indirect exposure to volatility because of their additional dependence on the price of the underlying. A volatility derivative on the other hand, is a derivative where the underlying instrument determining the payoff of the contract is some measure of volatility, which gives direct exposure to the volatility of the asset. Loosely speaking, it gives an investor a way to directly speculate on or hedge against the "riskiness" of an asset. Additionally, it is considered a stylized fact that the volatility of an equity asset generally increases as the price falls, i.e. there is a negative correlation between equity's price and volatility.

Thus buying volatility provides a crude hedge for a portfolio which is net long one or many equity assets.

In the 1990's, several volatility indices were born, and toward the end of that decade volatility derivatives became increasingly common [5]. The CBOE's volatility index (VIX), which is perhaps the most popular volatility index, was introduced in 1993 but then overhauled in 2003 to reflect the expected volatility of the S&P 500 over the coming month[1]. VIX futures were offered for the first time followed by VIX options in 2006. The CBOE touts the strong negative correlation between the VIX and the S&P 500 as motivation for trading volatility derivatives [6]. Volatility derivatives have become increasingly popular in recent years. According to data provided by the CBOE, the average daily trading volume of VIX options reached a new high of 792,668 contracts in January 2014 representing a 40% increase from 2013 volumes which were in turn 28% above 2012 volumes. Despite the growing interest, there does not yet seem to be a consensus on a pricing model or methodology.

The focus of this study is to examine four basic types of volatility derivatives: variance swaps, volatility swaps, and standard European-style call options on variance and volatility, concluding with method for pricing VIX options. In the following sections we will examine the ingredients of volatility derivatives as well as the structure of the four basic volatility derivatives themselves.

1.2 Measures of volatility

The most important ingredient in a volatility derivative is the measure of volatility. Most volatility derivatives are contracts based on the realized volatility, denoted by $R_{\tau,T}$, of a process over some period of time $[\tau, T]$, which can be described as discretely or continuously monitored. Practically, continuous monitoring is not possible so all volatility derivatives which are contracts on realized volatility are based on a discretely monitored measure of realized volatility, which is usually annualized. The monitored process is usually the prices of an asset or an index, but in general any measurable process could be monitored. If an asset price is modeled as an Itô process, then continuously monitored variance of the process described by the stochastic differential equation $dS_t = \mu_t S_t dt + \sigma_t S_t dW_t$ is the quadratic variation of the process $d\log S_t$, the square root of which is the realized volatility. Taking $\tau = t_1 < t_2 < ... < t_n = T$,

[1]The VIX actually is formulated to provide the square root of the fair strike price of a variance swap expiring in one month.

$$R^2_{\tau,T} := \frac{1}{T}\sum_{i=1}^{n}\log^2\frac{S_{t_i}}{S_{t_{i-1}}} \xrightarrow{p} \frac{1}{T-\tau}\int_{\tau}^{T}\sigma_t^2 dt, \quad as\ n \to \infty \tag{1.1}$$

Thus, as the partition size decreases, the discretization converges in probability to the realized variance of the process S_t (see appendix A.1). In practice, a calculation of realized volatility generally involves sampling the price of the associated asset at the end of each business day, resulting in the following typical calculation of realized volatility [5].

Example calculation of realized volatility (annualized) in basis points.

- $Y = 252$ - The number of trading days in a year
- n - the number of trading days in the period $[\tau, T]$
- S_{t_i}- the asset price at the end of day i
- $u^2 = Y/n \times 100^2$ - the scale factor, which in this case, converts R to an annualized percentage
- $R_{\tau,T} = \left(u^2 \sum_{i=1}^{n} \log^2 \frac{S_{t_i}}{S_{t_{i-1}}}\right)^{\frac{1}{2}}$

1.3 Variance swaps

The variance swap is the most tenable volatility derivative. Replication and hedging can be accomplished with a properly weighted strip of options, assuming only integrability of the volatility process making pricing essentially model free. For this reason, variance swaps initially enjoyed the most success out of the volatility derivatives with many banks offering these OTC contracts on both indices and individual stocks beginning in the late 1990's [5]. The variance swap has three ingredients:

- K^2_{vol} the fixed leg, or strike, usually quoted in volatility
- $R^2_{\tau,T}$ the floating leg, or realized variance
- N_{vega}, N_{var} the notional in terms of vega and variance, respectively

Combined, the payoff at time T is

$$N_{var}\left(R^2_{\tau,T} - K^2_{vol}\right) \tag{1.2}$$

where the strike is chosen such that the initial value of the swap is zero, i.e. the strike is equal to the expected realized variance over the lifetime of the swap. It may not be immediately clear where the vega notional comes into play or why the swap would be quoted in terms of volatility. From a practitioners point of view, it is desirable that the contract pays a fixed amount per point of realized volatility above/below the strike; this is the vega notional. A contract which pays a linear function of vega notional times realized volatility is a volatility swap (which will be covered later), but in order to make the variance swap more intuitive, the variance notional is chosen so that the variance swap approximates a volatility swap with vega notional and strike K_{vol}. More precisely, the derivative of a volatility swap with respect to volatility is vega notional. Taking the derivative of the variance swap with respect to volatility, evaluated at the volatility strike and setting it equal to the vega notional yields the following:

$$N_{vega} = \frac{d}{dR_{\tau,T}} N_{var} \left(R_{\tau,T}^2 - K_{vol}^2 \right) |_{R_{\tau,T}=K_{vol}} = 2N_{var}K_{vol} \tag{1.3}$$

Since the vega notional is given, the variance notional can be computed so that the variance swap pays approximately vega notional for each volatility point above the strike, but as the realized variance moves farther from the strike this approximation becomes increasingly worse. Perhaps these conventions can be made clearer by the following example, based on the handbook provided by Biscamp and Weithers [7].

Example variance swap:

- Contract: the buyer shall receive/pay approximately the vega notional, $N_{vega} = \$100,000$, for every point of realized volatility above/below the strike K_{vol} at the date of expiration T.
- Strike: the variance swap is struck such that its initial value is zero. For example, we take $\mathbb{E}\left[R_{0,T}^2\right] = K_{vol}^2 = 16^2$, annualized.
- Variance notional: the variance notional is chosen so that the buyer receives approximately the vega notional

$$N_{var} = \frac{N_{vega}}{2K_{vol}} = \frac{100,000}{2 \times 16} = 3,125$$

- Contract function: the exact payout of the contract is given by 1.2.
- Hypothetical payout at time T:
  - Scenario 1: the realized variance is 17^2, payout: $3,125\left(17^2 - 16^2\right) = 103,125$
  - Scenario 2: the realized variance is 15^2, payout: $3,125\left(15^2 - 16^2\right) = -96,875$

In the example above, we are thinking about volatility but trading variance, even though the payout of a variance swap is nonlinear in volatility. If we wanted linear exposure to volatility, perhaps trading a volatility swap would make more sense.

1.3.1 Variance swaps and the VIX

The VIX, as described earlier, was introduced by the CBOE to provide an indication of the level of expected volatility of the S&P 500 over the coming 30 days. The calculation which produces the current value of the VIX, however, is based on the fair strike of a variance swap starting today and expiring in 30 days ($T = 30$), more precisely, it is the square root of the expectation of realized variance.

$$VIX = \sqrt{\mathbb{E}\left[R_{0,T}^2\right]} \tag{1.4}$$

The exact calculation is described in the VIX white paper [6], which specifies how the options from the current month and the following month are weighted in order to produce a nearly continuous estimate. The rationale for the weighting scheme will be discussed in more depth when variance swap pricing is addressed in later sections. It is important to note, however, that the VIX is not an estimate of expected volatility but rather the square root of expected variance and the two are not equal. By an application of Jensen's inequality to a concave function we see that $\sqrt{\mathbb{E}\left[R_{0,T}^2\right]} \geq \mathbb{E}\left[R_{0,T}\right]$ is all we can say in terms of relating the expected realized volatility to the VIX.

1.4 Volatility swaps

A volatility swap is perhaps the most natural choice of volatility derivative for an individual wishing to hedge a portfolio against volatility exposure. Its structure is very similar to that of a variance swap.

- K_{vol} the fixed leg, or strike in volatility points
- $R_{\tau,T}$ the floating leg, or realized volatility
- N_{vega} the notional in terms of vega (currency per volatility point)

Combined, the payoff at time T is

$$N_{vega}\left(R_{\tau,T} - K_{vol}\right) \tag{1.5}$$

Since the payoff function of a volatility swap is both simple and intuitive for a trader, an example is not provided. The buyer of such a swap simply receives/pays the vega

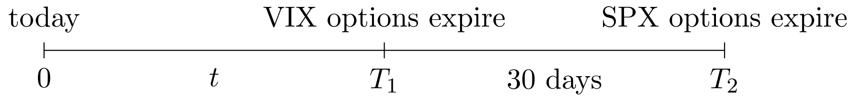

FIGURE 1.1: The VIX time line

notional for each point of realized volatility above/below the strike at the settlement date. The trouble with a volatility swap is that estimating future realized volatility is not straightforward and replicating a volatility is even more difficult, making pricing the swap during its lifetime much more cumbersome than pricing a variance swap. This difference will become obvious in later sections, but an intuitive explanation is that expectations of realized variance can be summed linearly in time whereas expectations of realized volatility cannot. In general, payoff functions that are linear in realized variance are easy to replicate and price, whereas nonlinear functions of realized variance, e.g. realized volatility, are not. Additionally, when pricing a volatility swap one has to choose between an approximate model free approach or applying some type of stochastic volatility model, where the swap price is determined within the model. All these considerations in tandem have made volatility swaps less popular than variance swaps despite the seemingly intuitive payoff function.

1.5 Options on variance and volatility

Standard European style options on variance and volatility are a natural extension of swaps. For example, a trader who is concerned only with hedging against an increase in volatility would prefer an option to a swap. The payoff functions at expiry T have the standard recognizable form:

- variance call: $\max\left(R^2_{\tau,T} - K^2_{vol}, 0\right) \quad or \quad \left(R^2_{\tau,T} - K^2_{vol}\right)^+$
- volatility call: $\max\left(R_{\tau,T} - K_{vol}, 0\right) \quad or \quad \left(R_{\tau,T} - K_{vol}\right)^+$

Put options are also formed in the standard way.

1.5.1 VIX options

Options on the VIX are actually options on an estimate of forward volatility of the S&P 500, and are therefore directly dependent on the SPX option prices. The time line in figure 1.1 relates the various dates necessary for understanding the structure of the VIX option.

Consider a European call option on the VIX expiring at T_1. The VIX at expiration will be calculated with options opening at $t = T_1$ and expiring at $T_2 = T_1 + 30$ days - precisely the third Friday of the following month - and we wish to know the value of the option at values of $t \in [0, T_1]$. We can then express the payoff of a VIX option with strike K_{VIX} in the following way:

$$(VIX_0 - K_{VIX}) = \left(\sqrt{\mathbb{E}\left[R^2_{T_1,T_2}\right]} - K_{VIX}\right)^+ \tag{1.6}$$

First, note that all VIX options contracts have a multiplier of $100. It is also important to note that VIX options expire on the Wednesday 30 days prior to the third Friday of the following month (time T_2), which is the day that SPX options opened at T_1 expire. This ensures that the VIX option payoff is dependent on options of only one expiry. These special considerations combined with unusual market dynamics, for example upward sloping Black implied volatility skew, make the VIX options notoriously difficult to price [8]. There are many different approaches to pricing VIX options. Wang and Diagler [8] evaluate four methods which broadly cover the spectrum of available approaches. The most basic approach, presented by Whaley [9], makes use of the Black-76 formula thereby imposing all the assumptions of the Black-Scholes model on VIX options. Grunbichler and Longstaff [10] model the VIX as a mean reverting process and price VIX options based on that process. Lin and Chang apply a stochastic volatility model with jumps in both the price and volatility process to the S&P 500 index to derive a VIX option pricing model. The Carr-Lee [11] approach is unique in that it attempts to price VIX options with the underlying S&P 500 options without applying any model to the S&P 500 price process. The result is a VIX option pricing formula which is a function of S&P 500 option prices and the VIX futures price. It is for this relative lack of model assumptions that the Carr-Lee method is attractive and evaluated in this paper.

Chapter 2

Pricing volatility derivatives

2.1 Replication of variance swaps

As discussed in the background, a variance swap is a contract, which at expiry T, pays a notional times the following

$$R^2_{0,T} - K_{var} \tag{2.1}$$

where $R^2_{0,T}$ is the realized variance over the period $[0, T]$ and K_{var} is the fixed leg of the swap. For this section, given $R^2_{t,T}$ we set $u^2 = 1/(T-t)$, meaning realized variance is not reported as a percentage but is annualized. The realized variance can be either continuously or discretely monitored, but in the following analysis we only consider the continuously monitored case.

Following the methods laid out in Neuberger [2] and expanded upon in Demeterfi, Derman et al [12], assume that the price process of a risky asset S_t is continuous, strictly positive, pays no dividends and satisfies the following SDE

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_t dW_t \tag{2.2}$$

where W is a Brownian motion under the objective probability measure, the drift μ and the volatility σ are assumed to be finite on the domain of interest and subject to some general conditions which ensure predictability (see appendix A.1). The volatility could be drawn from a predefined volatility surface, for example. Now examine the SDE for the process $df(t, S_t) = d(\log S_t)$. Itô 's lemma can be applied noting that the partial derivatives of f evaluated at (t, s) are $f_t = 0$, $f_s = 1/s$, $f_{ss} = -1/s^2$. Then

$$d(\log S_t) = 0 + \left(\frac{\mu_t S_t}{S_t} - \frac{\sigma_t^2 S_t^2}{2S_t^2}\right) dt + \frac{\sigma_t S_t}{S_t} dW_t = \left(\mu_t - \frac{\sigma_t^2}{2}\right) dt + \sigma_t dW_t \tag{2.3}$$

We can then isolate the variance by taking the difference between $dS_t/S_t$and $d(\log S_t)$ as follows:

$$\frac{dS_t}{S_t} - d(\log S_t) = \frac{\sigma_t^2}{2}dt \tag{2.4}$$

Integrating provides the following formula for replicating the continuously monitored realized variance $R_{0,T}^2$ over the time interval $[0, T]$.

$$R_{0,T}^2 = \frac{1}{T}\int_0^T \sigma_t^2 dt = \frac{2}{T}\int_0^T \frac{dS_t}{S_t} - \frac{2}{T}\log\frac{S_T}{S_0} \tag{2.5}$$

The first term on the right hand side represents the cumulative value of a dynamic position which holds $1/S_t$ shares of S_t at each time $t \in [0, T]$, that is, it is a position continuously rebalanced to be worth unity, the risk-neutral expectation of which is $2r$. The second term represents a static short position in a contract paying log of the return over the period. We wish to replicate the realized variance using instruments available in the market and log contracts are not traded, so a surrogate is needed for the payoff of the of the log contract.

Using the method prescribed in Carr-Madan [4], the payoff from any continuous, twice-differentiable contract function f can be replicated by

$$f(S_T) = f(\kappa)+f'(\kappa)(S_T-\kappa)+\int_0^\kappa f''(K)(K-S_T)^+dK+\int_\kappa^\infty f''(K)(S_T-K)^+dK \tag{2.6}$$

where κ is an arbitrary (positive) separator. See Appendix A.2 for a simple proof of this relation. Taking $f(s) = -\log s$ with derivatives $f'(s) = -1/s\, f''(s) = 1/s^2$ we arrive at the following replication of the payoff of being short one log contract.

$$-\log S_T = -\log\kappa - \frac{S_T}{\kappa} + 1 + \int_0^\kappa \frac{(K-S_T)^+}{K^2}dK + \int_\kappa^\infty \frac{(S_T-K)^+}{K^2}dK \tag{2.7}$$

Equation 2.7 tells us that the payoff from $-\log S_T$ can be replicated by a combination of

- Short$\frac{1}{\kappa}$ forward contracts on the stock
- A cash position of $-\log\kappa + 1$
- dK/K^2 European calls struck at $K > \kappa$and puts struck at $K < \kappa$

Define the position in a continuum of options as

$$I_{PC}(t,\kappa) := \int_0^\kappa \frac{P_t(K)}{K^2}dK + \int_\kappa^\infty \frac{C_t(K)}{K^2}dK \tag{2.8}$$

where $P_t(K)$ and $C_t(K)$ are the prices of the put and call options at time t struck at K. Incorporating the replication of the log contract into 2.5, we find

$$R_{0,T}^2 = \frac{2}{T}\left(\int_0^T \frac{dS_t}{S_t} + \log\frac{S_0}{\kappa} - \frac{S_T}{\kappa} + 1 + \int_0^\kappa \frac{(K-S_T)^+}{K^2}dK + \int_\kappa^\infty \frac{(S_T-K)^+}{K^2}dK\right) \tag{2.9}$$

As most variance swaps are struck such that they have an initial value of zero, it is interesting to know the expected value of the realized variance as this will be the swap rate K_{var}. Assuming that the risk-free rate is constant, the risk neutral expectation[1] of the realized variance over $[0,T]$ is

$$K_{var} = \mathbb{E}\left[R_{0,T}^2\right] = \frac{2}{T}\left(rT + \log\frac{S_0}{\kappa} - \frac{S_0e^{rT}}{\kappa} + 1 + e^{rT}I_{PC}(0,\kappa)\right) \tag{2.10}$$

Since κ is arbitrary we can set it equal to the initial forward price $F_0 := S_0e^{rT}$. In this case 2.10 simplifies to

$$K_{var} = \mathbb{E}\left[R_{0,T}^2\right] = \frac{2e^{rT}}{T}I_{PC}(0,F_0) \tag{2.11}$$

In summary, by assuming that the price process is positive and continuous, there is no arbitrage and the risk-free rate is constant we have developed a method for determining the expected value of the realized variance -the swap rate- based on traded instruments; specifically, holding a portfolio containing a weighted strip of put options below and call options above the forward price.

2.2 Pricing realized variance

Following the approach of Carr-Lee [13], we define the synthetic variance swap to be a self-financing portfolio which replicates $R_{0,T}^2$ where A_t is the present value of the expected payoff of a synthetic volatility swap. In other words, A_t is the value of a portfolio which replicates a variance swap with a fixed leg of zero. The value of A_t can be inferred from 2.10 given the information generated by the process S_t on the interval $[0,t]$.

$$\begin{aligned} A_t &= e^{-r(T-t)}\mathbb{E}_t\left[R_{0,t}^2 + R_{t,T}^2\right] && (2.12)\\ &= e^{-r(T-t)}\left(R_{0,t}^2 + \frac{2}{T-t}\left(r(T-t) + \log\frac{S_t}{\kappa} - \frac{S_te^{r(T-t)}}{\kappa} + 1\right)\right) + \frac{2}{T-t}I_{PC}(t,\kappa) \end{aligned}$$

[1] $\mathbb{E}_t$ denotes the expectation taken with respect to the risk neutral probability measure conditional on $\mathcal{H}_t$ (as defined in A.1) unless stated otherwise.

Defining the forward price to be $F_t = e^{r(T-t)}S_t$, the above simplifies to

$$A_t = e^{-r(T-t)}\left(R_{0,t}^2 + \frac{2}{T-t}\left(\log\frac{F_t}{\kappa} - \frac{F_t}{\kappa} + 1\right)\right) + \frac{2}{T-t}I_{PC}(t,\kappa) \tag{2.13}$$

Examining the synthetic variance swap portfolio, we can see it consists of a static position in options entered at $t = 0$, expiring at T, and a dynamic cash account, specifically

- $e^{-r(T-t)}\left(R_{0,t}^2 + \frac{2}{T-t}\left(\log\frac{F_t}{\kappa} - \frac{F_t}{\kappa} + 1\right)\right)$ cash (dynamic)
- $\frac{2}{(T-t)K^2}dK$ European calls struck at $K > \kappa$and puts struck at $K < \kappa$ (static)

In order to write this in the form of Carr-Lee [13], we note that

$$e^{-r(T-t)}\left(1 - \frac{F_t}{\kappa}\right) = \left(\frac{1}{F_t} - \frac{1}{\kappa}\right)S_t$$

and the contents of the synthetic variance swap portfolio can be expressed as

- $e^{-r(T-t)}\left(R_{0,t}^2 + \frac{2}{T-t}\log\frac{F_t}{\kappa}\right)$ cash (dynamic)
- $\left(\frac{1}{F_t} - \frac{1}{\kappa}\right)$ shares (dynamic)
- $\frac{2}{(T-t)K^2}dK$ European calls struck at $K > \kappa$ and puts struck at $K < \kappa$ (static)

Again following Carr-Lee [13], the expectation at $t = 0$ of the payout of a synthetic variance swap as a function of the ratio of the stock price at expiry to the current forward price is plotted in figure 2.1. At $t > 0$, as variance accumulates, the curve simply shifts upwards.

2.3 Volatility swaps: pricing and replication

Pricing and replication of volatility swaps presents a new set of difficulties. Realized volatility has a nonlinear time-dependence and leverage effects - or the correlation between volatility and the associated price process - cannot be ignored. Below, we examine a well known method for pricing and replicating a volatility swap which allows for correlation between the volatility and price processes and imposes minimal constraints on the volatility process. The performance of this method is evaluated using the Heston model and compared to an estimate of volatility swap value approximated by the forward-at-the-money Black-Scholes implied volatility under a zero correlation assumption.

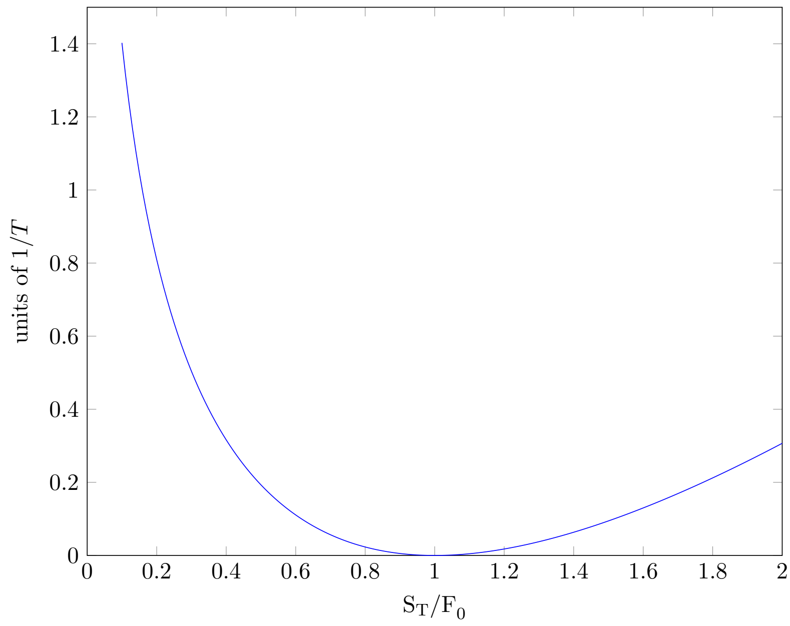

FIGURE 2.1: The expected payoff of a synthetic variance swap

2.3.1 The Carr-Lee method for replicating volatility swaps

Using the methods introduced by Carr-Lee [11], which assume the same volatility/price correlation behavior used in many stochastic volatility models such as the Heston model, a volatility swap can be both priced and replicated approximately by a dynamic position in European options on the underlying. The method is exact for zero correlation only, however it approximates the value of a volatility swap under non-zero correlation much better than a simple estimate based on Black-Scholes implied volatility.

The Carr-Lee approach has four main components:

1. The mixing formula, which describes the underlying price process and correlation with the volatility process.
2. The correlation-neutral condition, which describes the condition necessary to price a correlation-neutral claim.
3. The correlation-neutral exponential claim, which is a first order invariant (w.r.t. correlation) exponential claim on realized variance, replicable by European options on the underlying.
4. The correlation-neutral volatility swap, which admits replication by correlation-neutral exponential claims.

Drawing directly from [11], the components listed above are summarized in the following paragraphs.

2.3.1.1 The mixing formula

The underlying is assumed to have the following dynamics, assuming denomination by an asset which pays 1 at time T:

$$dS_t = \sigma_t S_t \left(\sqrt{1-\rho^2} dW_{1t} + \rho dW_{2t} \right) \tag{2.14}$$

where $|\rho| \leq 1$ is the correlation between the volatility and price processes, W_1 and W_2 are independent$\mathcal{F}_t$-Brownian motions, σ_t is the volatility process which is adapted to some filtration $\mathcal{H}_t \subseteq \mathcal{F}_t$ and independent of $\mathcal{F}_t^{W_1}$. This is identical to the correlation assumption made by most stochastic volatility models, here however, no further model assumptions are imposed on the volatility process.

2.3.1.2 The correlation-neutral condition

Define G to be a function which maps S_T to a European style payoff, e.g. $G(S_T) = (S_T - K)^+$and let $G^{BS}(S_t, \sigma)$ be the expectation of$G$ per the Black-Scholes formula. Then by a Taylor expansion of $\mathbb{E}_t G(S_T)$ about $\rho = 0$, the following first order approximation can be made:

$$\mathbb{E}_t G(S_T) \approx \mathbb{E}_t G^{BS}(S_t, \bar{\sigma}) + O(\rho^2) \tag{2.15}$$

if "the contract's Black-Scholes delta is constant across all volatility parameters" (Carr-Lee, 11), specifically:

$$\frac{\partial G^{BS}}{\partial S_t}(S_t, \sigma) = const. \quad \text{for all constants } \sigma \geq 0 \tag{2.16}$$

where $\bar{\sigma}^2 := \int_t^T \sigma_s^2 ds$. By enforcing the condition above, the claim described by F will be first order correlation-neutral.

2.3.1.3 The correlation-neutral exponential claim

The correlation neutral exponential claim is found by pricing correlation-sensitive exponential claims on realized variance, then exploiting the non-uniqueness of the solutions to enforce 2.16. Expressing the realized variance as the quadratic variation of the log price process $\langle X \rangle_t$, the correlation-neutral exponential claim is replicated by:

$$\mathbb{E}_t e^{\lambda \langle X \rangle_T} = e^{\lambda \langle X \rangle_t} \mathbb{E}_t \left[(1-\theta) \left(S_T/S_t \right)^p + \theta \left(S_T/S_t \right)^{1-p} \right] \tag{2.17}$$

where

$$\theta(\lambda) = \frac{1}{2} + \frac{1}{2\sqrt{1+8\lambda}} \tag{2.18}$$

$$p(\lambda) = \frac{1}{2} + \frac{1}{2}\sqrt{1+8\lambda} \tag{2.19}$$

The function $\theta(\lambda)$ was chosen specifically from a family of functions in order to ensure correlation-neutrality as per 2.16.

2.3.1.4 Correlation-neutral replication of volatility swaps

In order to relate the expectation of realized volatility with the correlation-neutral exponential claim, Carr-Lee make use of a well known formula, which, applied to realized variance takes the following form:

$$\mathbb{E}_t\sqrt{\langle X\rangle_T - \langle X\rangle_t} = \frac{1}{2\sqrt{\pi}}\int_0^\infty \frac{1-\mathbb{E}_t \exp\left\{-\lambda\left(\langle X\rangle_T - \langle X\rangle_t\right)\right\}}{\lambda^{3/2}} d\lambda \tag{2.20}$$

Applying the correlation-neutral exponential yields

$$\mathbb{E}_t\sqrt{\langle X\rangle_T} = \frac{1}{2\sqrt{\pi}}\mathbb{E}_t\int_0^\infty \frac{1+(\theta-1)e^{-\lambda\langle X\rangle_t}\left(\frac{S_T}{S_t}\right)^p - \theta e^{-\lambda\langle X\rangle_t}\left(\frac{S_T}{S_t}\right)^{1-p}}{\lambda^{3/2}} d\lambda \tag{2.21}$$

The power claims in the integrand can be replicated using European options expiring at T via 2.6. Carr-Lee then integrate these claims over λ to obtain the replicating portfolio for $\mathbb{E}_0\sqrt{\langle X\rangle_T}$ (K denotes the strike of a European option expiring at T and F is the T-forward price of the underlying).

- $\sqrt{\pi/2}/F_0$ straddles at $K = F_0$
- $\sqrt{\frac{\pi}{8K^3F_0}}\left[I_1\left(\frac{1}{2}\log\frac{K}{F_0}\right) - I_0\left(\frac{1}{2}\log\frac{K}{F_0}\right)\right] dK$ calls at strikes $K > F_0$
- $\sqrt{\frac{\pi}{8K^3F_0}}\left[I_0\left(\frac{1}{2}\log\frac{K}{F_0}\right) - I_1\left(\frac{1}{2}\log\frac{K}{F_0}\right)\right] dK$ puts at strikes $K < F_0$

Where I_ν denotes the modified Bessel function of order ν. For times $t \in (0,T)$ when $\langle X\rangle_t > 0$, the net position replicating $\mathbb{E}_t\sqrt{\langle X\rangle_T}$ is

- $\frac{dK}{\sqrt{\pi}}\int_0^\infty \frac{e^{-\lambda\langle X\rangle_t}}{K^2\sqrt{\lambda}}\left[(1-\theta)\left(K/F_t\right)^p + \theta\left(K/F_t\right)^{1-p}\right] d\lambda$ calls at $K > F_0$, puts at $K < F_0$
- $\sqrt{\langle X\rangle_t}$ bonds

In order to express this value as an annualized amount we divide by $\sqrt{T-t}$. We will refer to this portfolio as the synthetic volatility swap.

$$B_t = \frac{1}{\sqrt{T-t}} \mathbb{E}_t \sqrt{\langle X \rangle_T}$$

2.3.2 Approximations under zero-correlation

The volatility swap rate, or the current fair value of expected realized volatility can be very closely approximated by the forward at the money (ATM) implied volatility of a European option expiring with the swap, *if* one assumes zero correlation between the price and volatility processes. Carr-Lee [5], show this result in detail giving credit to Feinstein [14] for originally publishing the method and Brenner-Subrahmanyam [15] for their approximation of the at ATM option value under Black-Scholes. Perhaps the most illuminating explanation though comes as a result of Friz & Gatheral's [16] formula for realized volatility under zero correlation:

$$\mathbb{E}_0 \sqrt{\langle X \rangle_T} = \frac{\sqrt{2\pi}}{F_0} C(F_0) + \sqrt{\frac{2}{\pi}} \int_{F_0}^{\infty} \sqrt{\frac{1}{K^3 F_0}} I_0 \left(\frac{1}{2} \log \frac{K}{F_0} \right) C(K) dK$$

The second term in the formula above accounts for only a very small percentage of the total expectation. As Brenner-Subrahmanyam showed, a Taylor expansion of the ATM Black-Scholes call formula implies

$$C^{BS} \approx \frac{F_0 \sigma^{BS} \sqrt{T}}{\sqrt{2\pi}}$$

Thus to very close approximation

$$\mathbb{E}_0 \sqrt{\langle X \rangle_T} \approx \sigma^{BS} \sqrt{T}$$

It is also interesting to note that the approximation error, $\sigma^{BS}\sqrt{T} - \mathbb{E}_0\sqrt{\langle X \rangle_T}$, has the same sign as the ATM skew of the volatility smile. That is, for negative correlation (negative skew) the difference is negative and vice versa, which becomes clear when examining figure 2.2.

2.3.3 Numerical simulations with Heston dynamics

Using the Heston model and assuming the parameters suggested by Bakshi, Cao and Chen [17] with a time to expiry $T = 0.5$ years, the performance of the Carr-Lee approach was evaluated. First, the expected value of realized volatility was calculated directly from the Heston volatility process (see appendix A.3), which is independent of correlation.

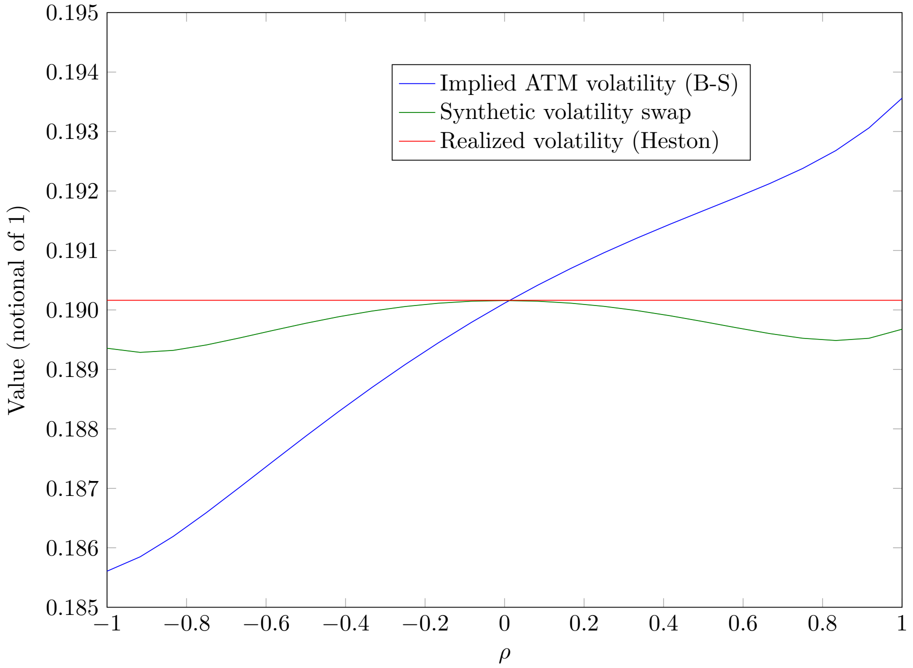

FIGURE 2.2: Synthetic volatility swap performance in a Heston model simulation

A strip of option prices were generated via the known Heston model pricing formulas for a range of correlation values $\rho \in [-1, 1]$ using the fractional FFT method described in [18]. These options were then used to construct a synthetic volatility swap portfolio unique to each correlation value. Figure 2.2 shows the value of the synthetic volatility swap portfolio as a function of correlation ρ, as well as the expected value of realized volatility inherent to the chosen Heston parameters and the Black-Scholes implied volatility. Note that the synthetic volatility swap value is exact and has zero slope for zero correlation. The synthetic volatility swap only approximates the value of the volatility swap at $\rho \neq 0$, but the approximation error is quite small; a maximum of 0.5% for $\rho \approx 0.9$ in this case. The approximation error increases with both time to expiry and volatility of volatility. If T is increased to 1, the maximum error increases to about 2% but this is for the unlikely case that $\rho = -1$. If the volatility of volatility is then increased to 0.5 the maximum error of 4.5% is again attained at $\rho = -1$. In rough terms, one could say that in a typical market the approximation error is capped at about 5% for expiries up to 1 year. For extreme market conditions the approximation erodes quickly, but the Heston model is not appropriate for these conditions either. In this case, it would be interesting to evaluate the Carr-Lee method in the context of a stochastic volatility model with jumps in the underlying.

As time passes and volatility accumulates, figure 2.3 shows how the expected payout evolves from the original position consisting primarily of a straddle, to a position which produces a payout very similar to that of a square root variance swap. This is because

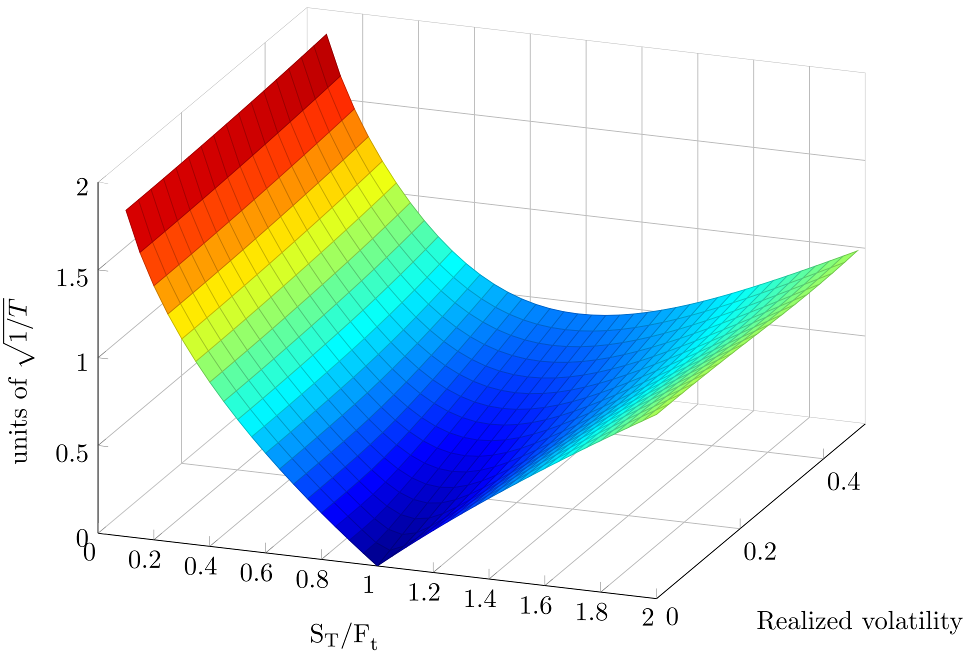

Figure 2.3: Evolution of the expected payout from a volatility swap

the concavity of the volatility swap payout decreases with increasing variance, making the expected payout of a volatility swap roughly equal to a variance swap plus cash, as Carr-Lee [13] note.

2.4 Options on realized volatility: a lognormal approach

Pricing options on realized variance and volatility can be made quite simple if one assumes, a priori, that the time-t conditional distribution of realized volatility is lognormal, as Friz and Gatheral[16] suggest. This approach is also followed by Carr-Lee[13]. The argument that Gatheral[19] makes in support of this assumption is essentially two-fold: "the empirical distribution of implied volatility changes looks lognormal" and "the dynamics of volatility skew are consistent with approximately lognormal volatility dynamics." Assuming that volatility dynamics are approximately lognormal, one would expect that dynamics of realized volatility to be approximately lognormal as well because the sum of lognormally distributed random variables is approximately lognormal. Adopting the lognormal assumption,

$$\log R_{t,T} \sim \mathcal{N}(\mu, s^2) \implies \log R^2_{t,T} \sim \mathcal{N}(2\mu, 4s^2)$$

Friz and Gatheral[16] deduce:

$$B_t := \mathbb{E}_t\left[R_{t,T}\right] = e^{\mu+s^2/2}, \quad A_t := \mathbb{E}_t\left[R_{t,T}^2\right] = e^{2\mu+2s^2} \tag{2.22}$$

Given that we know how to compute the expected value of both realized variance and volatility, we can solve for μ and s at t yielding the following, while suppressing the dependence on t.

$$\mu = 2\log B - \frac{1}{2}\log A \tag{2.23}$$

$$s^2 = \log A - 2\log B \tag{2.24}$$

We now have a lognormal distribution with known mean and variance. Thus, pricing variance and volatility options follows via the standard Black-Scholes style approach. First we examine a European call on realized variance over the period $[0, T]$ of which we wish to know the initial value.

$$C_0^{Var} = e^{-rT}\mathbb{E}_0\left[\left(R_{0,T}^2 - K\right)^+\right] \tag{2.25}$$

As per our assumption of lognormality,

$$\mathbb{P}\left[R_{0,T}^2 > K\right] = \mathbb{P}\left[e^{2\mu+2sz} > K\right] \tag{2.26}$$

where $Z \sim \mathcal{N}(0,1)$. Defining

$$d = \frac{\mu - \frac{1}{2}\log K}{s} \tag{2.27}$$

we can infer that the probability that the option expires in the money implies the following:

$$\mathbb{P}\left[R_{0,T}^2 > K\right] \Longleftrightarrow \mathbb{P}\left[Z > -d\right] \tag{2.28}$$

We can now take the expectation in 2.25 given the probability density and cumulative distribution functions of the standard normal

$$\varphi(z) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}z^2}, \qquad N(x) = \int_{-\infty}^{x}\varphi(z)dz$$

such that

$$\begin{aligned}
C_0^{Var}e^{rT} &= \int_{-d}^{\infty}\left(e^{2\mu+2sz}-K\right)\varphi(z)dz \\
&= \frac{1}{\sqrt{2\pi}}\int_{-d}^{\infty}e^{2\mu+2sz-\frac{1}{2}z^2}dz-KN(d) \\
&= \frac{1}{\sqrt{2\pi}}\int_{-d}^{\infty}e^{2\mu+2s^2-\frac{1}{2}(z-2s)^2}dz-KN(d) \\
&= e^{2\mu+2s^2}\int_{-\infty}^{d+2s}\varphi(z)dz-KN(d) \\
&= e^{2\mu+2s^2}N(d+2s)-KN(d)
\end{aligned}$$

which yields the initial value of a variance call we desired, expressed more simply as:

$$C_0^{Var}=e^{-rT}\left[A_0N(d+2s)-KN(d)\right] \tag{2.29}$$

Equation 2.29 can be easily extended to value the option at $t\in[0,T]$ by recognizing that realized variance is additive, resulting in

$$d=\frac{\mu-\frac{1}{2}\log\left(K-R_{0,t}^2\right)}{s} \tag{2.30}$$

for $K>R_{0,t}^2$, noting that μ and s^2 are functions of A_t and B_t. Then

$$C_t^{Var}=e^{-r(T-t)}\begin{cases}A_tN(d+2s)-\left(K-R_{0,t}^2\right)N(d) & for\, K>R_{0,t}^2\\ A_t+R_{0,t}^2-K & for\, K\leq R_{0,t}^2\end{cases} \tag{2.31}$$

The same technique can be applied to an option on realized volatility. Consider a European call on realized volatility over the period $[0,T]$ of which we wish to know the initial value.

$$C_0^{Vol}=e^{-rT}\mathbb{E}_0\left[(R_{0,T}-K)^+\right] \tag{2.32}$$

Defining

$$d=\frac{\mu-\log K}{s} \tag{2.33}$$

we can infer that the probability that the option expires in the money implies the following:

$$\mathbb{P}\left[R_{0,T}>K\right]\Longleftrightarrow\mathbb{P}\left[Z>-d\right] \tag{2.34}$$

where z is drawn from the standard normal distribution, then

$$\begin{aligned} C_0^{Vol} e^{rT} &= \int_{-d}^{\infty} \left(e^{\mu+sz} - K\right) \varphi(z) dz \\ &= \frac{1}{\sqrt{2\pi}} \int_{-d}^{\infty} e^{\mu+sz-\frac{1}{2}z^2} dz - KN(d) \\ &= \frac{1}{\sqrt{2\pi}} \int_{-d}^{\infty} e^{\mu+\frac{1}{2}s^2-\frac{1}{2}(z-s)^2} dz - KN(d) \\ &= e^{\mu+\frac{1}{2}s^2} N(d+s) - KN(d) \end{aligned}$$

Which yields the initial value of a volatility call:

$$C_0^{Vol} = e^{-rT} \left[B_0 N(d+s) - KN(d)\right] \tag{2.35}$$

Extending this formula to $t \in [0, T]$ requires numerical integration, however. Now we wish to value

$$C_t^{Vol} = e^{-r(T-t)} \mathbb{E}_t \left[(R_{0,T} - K)^+\right] \tag{2.36}$$

$$= e^{-r(T-t)} \mathbb{E}_t \left[\left(\sqrt{R_{0,t}^2 + R_{t,T}^2} - K\right)^+\right] \tag{2.37}$$

which, following the same methods, yields

$$d = \frac{\mu - \frac{1}{2}\log\left(K^2 - R_{0,t}^2\right)}{s} \tag{2.38}$$

$$C_t^{Vol} = e^{-r(T-t)} \begin{cases} \int_{-d}^{\infty} \left(e^{2\mu+2sz} + R_{0,t}^2\right)^{\frac{1}{2}} \varphi(z) dz - KN(d) & for\, K > R_{0,t} \\ \mathbb{E}_t \left[R_{0,T}^2\right] - K & for\, K \leq R_{0,t} \end{cases} \tag{2.39}$$

2.4.1 VIX options and future realized volatility

Recall the structure of VIX options from section 1.5.1 of the background; we wish to price an option on the square root of the expected value of realized variance. A call expiring at T_1 has the payoff

$$(VIX_{T_1} - K)^+ = \left(\frac{100}{\sqrt{T_1 - T_1}} \sqrt{\mathbb{E}\left[R_{T_1,T_2}^2\right]} - K\right)^+$$

where R is the unannualized, unscaled realized volatility; i.e. scale factor $u = 1$. Throughout this section, all scaling due to the VIX index being reported in annual basis points will be stated explicitly. The associated time line is shown in 2.4 again

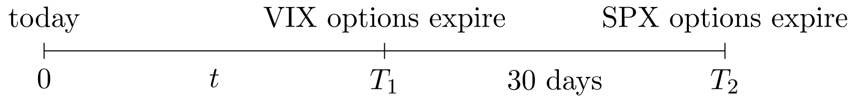

FIGURE 2.4: The VIX time line

for reference. As before, we assume that realized volatility is lognormal. Additionally, we assume that R^2_{0,T_1} and $R^2_{T_1,T_2}$ are independent. In order to price VIX options, it is necessary to compute the expectation of realized volatility and variance as well as the corresponding distribution parameters over the three following periods: $[0, T_1]$, $[0, T_2]$ and $[T_1, T_2]$, where the corresponding parameters are (μ_1, s_1^2), (μ_2, s_2^2) and (μ_{12}, s_{12}^2) respectively. For example, the distribution parameters for the period $[0, T_1]$ are given by solving

$$\mathbb{E}_0\left[R_{0,T_1}\right] = e^{\mu_1 + s_1^2/2}, \quad \mathbb{E}_0\left[R^2_{0,T_1}\right] = e^{2\mu_1 + 2s_1^2} \tag{2.40}$$

The parameters (μ_2, s_2^2) and (μ_{12}, s_{12}^2) are computed in an identical manner from the corresponding expectations of variance and volatility over periods $[0, T_2]$ and $[T_1, T_2]$. The expectation of realized variance,$\mathbb{E}_0\left[R^2_{0,T}\right]$, can be computed from 2.10 and the expectation of realized volatility $\mathbb{E}_0\left[R_{0,T}\right]$ can be calculated via the method described in section 2.3.1.4. Thus it remains to show how we compute the expectation of future realized variance and volatility.

Proposition 1. *Given European options on the underlying expiring at T_1 and T_2, the expectation of future realized variance is*

$$\mathbb{E}_0\left[R^2_{T_1,T_2}\right] = 2\left(e^{rT_2} I_{PC}^{T_2}(t, S_0 e^{rT_2}) - e^{rT_1} I_{PC}^{T_1}(t, S_0 e^{rT_1})\right) \tag{2.41}$$

By definition, realized variance is additive in time so $R^2_{0,T_1} + R^2_{T_1,T_2} = R^2_{0,T_2}$. Thus, $R^2_{T_1,T_2}$ is straightforward to compute using equation 2.10, given options prices on the SPX expiring at T_1 and T_2. Unfortunately, there is no analogous expression for computing the expectation of forward realized volatility. However, we propose the following:

Proposition 2. *Assuming that R^2_{0,T_1} and $R^2_{T_1,T_2}$ are independent and lognormally distributed, the parameters (μ_{12}, s_{12}^2) are given by*

$$s_{12}^2 = \frac{1}{4}\log\left(\frac{Var\left(R^2_{T_1,T_2}\right)}{\mathbb{E}_t\left[R^2_{T_1,T_2}\right]^2} + 1\right) \tag{2.42}$$

$$\mu_{12} = \frac{1}{2}\log\left(\mathbb{E}\left[R^2_{T_1,T_2}\right]\right) - s_{12}^2 \tag{2.43}$$

where

$$Var\left(R^2_{T_1,T_2}\right) = Var\left(R^2_{0,T_2}\right) - Var\left(R^2_{0,T_1}\right) \tag{2.44}$$

if $Cov\left(R^2_{0,T_1}, R^2_{T_1,T_2}\right) = 0$ *or equivalently* $Cov\left(R^2_{0,T_1}, R^2_{0,T_2}\right) = Var\left(R^2_{0,T_1}\right)$.

Equations 2.42 and 2.43 are derived from expressing the mean and variance of $R^2_{T_1,T_2}$ in terms of μ_{12} and s^2_{12}. The variances $Var\left(R^2_{0,T_2}\right)$ and $Var\left(R^2_{0,T_1}\right)$ are computed from expectations of realized variance and volatility over the respective periods. At this point we have defined all the parameters for the lognormal distributions of expected variance and volatility over the three time periods in question.

2.4.2 VIX Futures

A VIX futures contract maturing at T_1, has the following price at maturity:

$$Y = \sqrt{\mathbb{E}_{T_1}\left[R^2_{T_1,T_2}\right]} \tag{2.45}$$

Which has a current value of

$$F_0^{VIX} = \mathbb{E}_0\sqrt{\mathbb{E}_{T_1}\left[R^2_{T_1,T_2}\right]} \tag{2.46}$$

That is, a VIX future is the current value of the expected root of future realized variance between T_1 and T_2. Through applications of Jensen's inequality, the following bounds can be placed on the futures price, such that

$$\mathbb{E}_0\left[R_{T_1,T_2}\right] \leq F_0^{VIX} \leq \sqrt{\mathbb{E}_0\left[R^2_{T_1,T_2}\right]} \tag{2.47}$$

The tightness of these bounds depends on the volatility of volatility; $\sqrt{\mathbb{E}_0\left[R^2_{T_1,T_2}\right]} - \mathbb{E}_0\left[R_{T_1,T_2}\right]$ increases as volatility of volatility is increases. Now examine the variance of of a VIX future:

$$Var_0\left[Y\right] = \mathbb{E}_0\left[Y^2\right] - \mathbb{E}_0\left[Y\right]^2 \tag{2.48}$$

Gatheral [19], who gives credit to Bruno Dupire for this insight, shows that a VIX future can be priced based on forward realized variance and its own variance, which he suggests simply estimating from historical prices.

$$F_0^{VIX} = \sqrt{\mathbb{E}_0\left[R^2_{T_1,T_2}\right] - Var_0\left[Y\right]} \tag{2.49}$$

2.4.3 VIX options prices via VIX futures

VIX options can priced using VIX futures if we assume that expected value of forward realized variance is lognormal, which is the solution that Carr-Lee [13] suggest. We set:

$$F_t^{VIX} := \mathbb{E}_t\sqrt{\mathbb{E}_{T_1}\left[R^2_{T_1,T_2}\right]} = e^{\mu+s^2/2} \tag{2.50}$$

$$A_t := \mathbb{E}_t\left[R^2_{T_1,T_2}\right] = e^{2\mu+2s^2} \tag{2.51}$$

$$B_t := \mathbb{E}_t\left[R_{T_1,T_2}\right] \tag{2.52}$$

where the VIX futures can be either taken at spot value or priced using 2.49 and it must hold that $\sqrt{A_t} > F_t^{VIX}$. Then, the same method applied to find the value of a call on a volatility swap can be applied here, yielding a slightly modified form of equation 2.39.

$$C_0^{VIX}\left(F_0^{VIX}, K\right) = e^{-rT}\left[F_0^{VIX}N(d+s) - KN(d)\right] \tag{2.53}$$

The downside to this approach is that we would prefer a method which prices VIX options based solely on the SPX options from which the VIX is computed. To the author's knowledge, there is no method that can accomplish this without making even more restrictive assumptions than we have here. However, using the prescribed bounds for VIX futures, we can prescribe an upper bound for a VIX call without relying on the market price of VIX futures.

Proposition 3. *If equation 2.53 holds, an upper bound for the price of a VIX call is given by*

$$C_0^{VIX}\left(F_0^{VIX}, K\right) \leq \max\left\{\sqrt{A_0} - K,\, C_0^{VIX}\left(B_0, K\right)\right\} \tag{2.54}$$

Support for this proposition is given in Appendix A, but an intuitive explanation is given here. As F increases/decreases, the variance s of the assumed lognormal distribution decreases/increases causing the curve $C(K)$ becomes more/less convex so that $C_0^{VIX}(B_0, K) > C_0^{VIX}\left(F_0^{VIX}, K\right)$ for K greater than their point of intersection and $\sqrt{A_0} - K > C_0^{VIX}\left(F_0^{VIX}, K\right)$ for K less than their point of intersection. These two points of intersection are nearly identical (the regions overlap slightly) resulting in the upper bound given in proposition 3. This upper bound is shown to hold numerically in Chapter 3.

Chapter 3

Model performance and market dynamics

The small market study presented here was primarily undertaken to address the viability of the model assumptions as well as examine practical considerations of implementation and comment on market behavior.

3.1 Data

The study examined VIX options expiring in April and May 2014 over the five weeks preceding the April expiry. The data, collected at closing from the CBOE's site on a weekly basis, consist of :

- full strips of VIX options (April and May expiry)
- full strips of SPX options (April, May and June expiry)
- VIX spot, SPX spot, VIX futures (April, May) prices
- 3-month Treasury Bill yield (from Bloomberg)

3.2 Method

VIX option prices were computed with the method proposed by Carr-Lee [13], given in equation 2.53 along with bounds. In order to compute the VIX option prices for the April expiry, the April VIX future was used and the expectation of forward volatility and variance were computed from SPX options expiring in April and May with the risk-free rate assumed to be the 3-month T-Bill yield. In order to be used in calculations, the option prices were required to have non-zero bid and ask prices as well as non-zero trading volume on the valuation day.

3.2.1 Technical obstacles

Two main technical obstacles were faced. The first obstacle was the lack of SPX options expiring with the VIX options. This was due to the fact that VIX options expire on a Wednesday thus the closest expiry of SPX options was April 19 for the VIX options expiring on April 16 and May 17 for the VIX options expiring on May 21. The second obstacle was that the options were sold at discrete strikes for a relatively small range around the SPX spot and the pricing method requires a continuum of strikes from zero to infinity. Both obstacles were mitigated by employing a parametrization of the Black-Scholes implied volatility surface as suggested by Gatheral [20]. By fitting Gatheral's SVI model to the implied volatilities, SPX option prices could be generated for any positive strike allowing the mesh to be refined. In order to generate prices of SPX options expiring with the VIX options, the implied volatility curve corresponding to the nearest SPX expiry was used. This could have been further improved by using a time dependent parametrization, but that improvement is left to future work.

When trading volume was low, problems with the data inclusion criteria arose. The June SPX option data collected on March 14 did not have sufficient trading volume to construct a reasonable implied volatility curve, so the May VIX option prices on that date could not be calculated. By changing the inclusion criteria and perhaps employing a suitable weighting scheme, it might be possible to use data with very sparse trading, but this was not investigated.

3.3 Results

The bounds on the VIX futures, displayed in tables 3.1 and 3.2, were calculated from 2.47 at each valuation date for both VIX expiries with the exception of May VIX futures at the March 14 valuation date.

	Mar 14	Mar 21	Mar 28	Apr 04	Apr 11
Upper bound	17.6574	16.3741	15.9424	14.4897	16.9709
Market	16.7700	15.7500	15.5600	14.4500	16.8400
Lower bound	14.8523	13.5355	13.2922	12.2868	14.7484

TABLE 3.1: April VIX futures: market prices and bounds

	Mar 21	Mar 28	Apr 04	Apr 11
Upper bound	17.1914	16.2161	15.7277	16.4480
Market	16.2000	16.2500	15.3700	16.6300
Lower bound	14.8032	14.7451	13.8188	14.9820

TABLE 3.2: May VIX futures: market prices and bounds

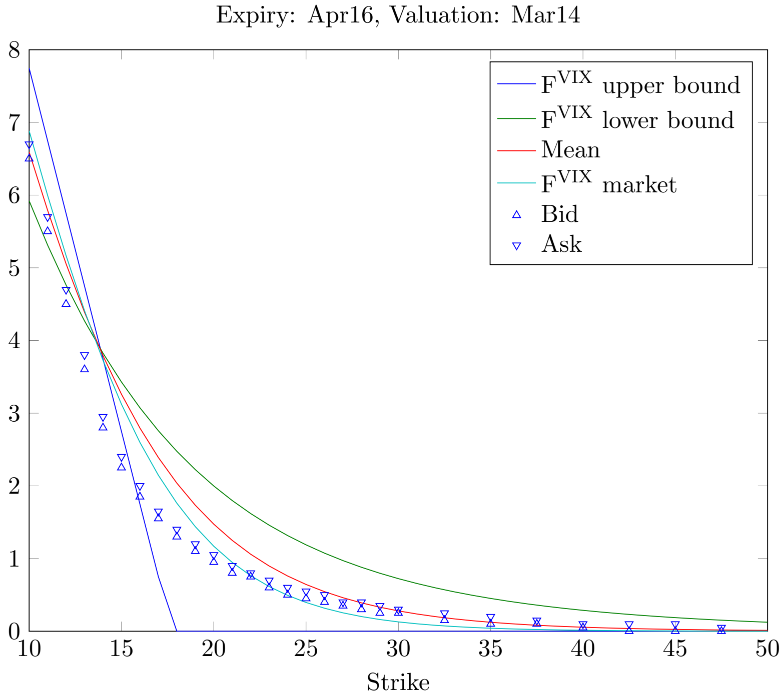

FIGURE 3.1: Comparison of theoretical and market VIX options prices

The call prices were calculated via equation 2.53 and plotted. Figure 3.1 depicts the results from the valuation of the April expiry on March 14; see Appendix B, figures B.2 and B.1 for all plots. The curve marked "F^{VIX} market" was generated by applying the formula proposed by Carr-Lee, given in equation 2.53. The curve marked "F^{VIX} upper bound" was generated by replacing the market price F^{VIX} with its upper bound in the valuation formula 2.53. Similarly, the curve marked "F^{VIX} lower bound" was generated by replacing F^{VIX} with its lower bound and the curve marked "mean" was generated by replacing F^{VIX} with the mean of the lower bound and upper bound.

The relative pricing error produced by equation 2.53 was computed and plotted in figures 3.2 and 3.3 as a function of forward moneyness for VIX options which had non-zero trading volume on the day of valuation. For the VIX options expiring in May, the relative error could only be computed on March 21 and April 4 for the following reasons: the March 14 prices could not be calculated due to lack of usable data, as discussed earlier, and the March 28 and April 11 option prices could not be calculated because the VIX futures price was above its theoretical upper bound which will be discussed further in the following section.

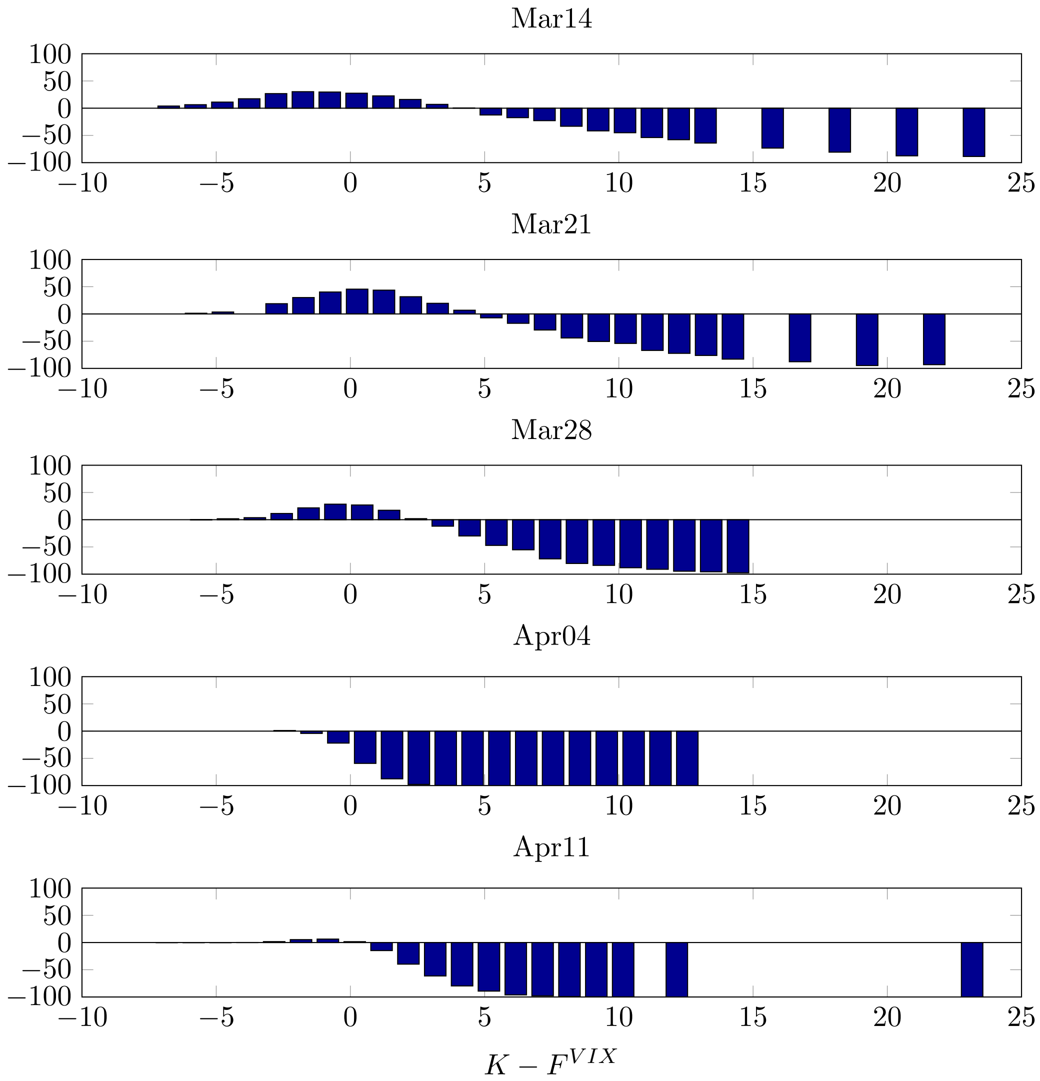

FIGURE 3.2: Percent difference between theoretical and market VIX options prices - April expiry

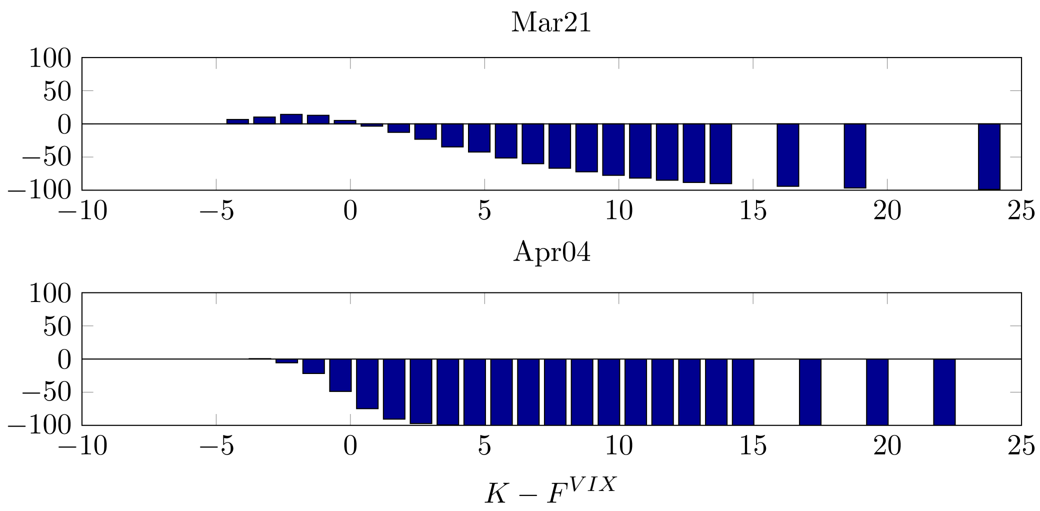

FIGURE 3.3: Percent difference between theoretical and market VIX options prices - May expiry

3.4 Discussion of results

Perhaps the most striking initial observation is that the May VIX futures price exceeded its upper bound on two occasions (March 28 and April 11) as can be seen in table 3.2. This is remarkable because the upper bound is essentially model free; the main limiting assumption being no jumps in the SPX. However, as Gatheral [19] suggests, the error induced by this assumption is on the order of the jump sizes cubed, i.e. very small. When the VIX futures price exceeds its upper bound, the Carr-Lee valuation formula cannot be used due to a violation of the underlying assumptions. This can be easily seen by setting $F^{VIX} = \sqrt{A}$ in equation 2.53, which results in the assumed lognormal distribution having zero variance. Thus for the dates where the VIX futures price exceeded its upper bound, the theoretical prices could not be computed which is reflected in figures 3.3 and B.2 depicting the errors and price curves respectively. Also, in those cases, the VIX options prices also exceed the upper bound as given by equation 2.54, a result unique to those cases. More generally, VIX futures traded close to their upper bound in all cases. In fact, for the period studied, the VIX futures traded not only above the lower bound but above the mean of the upper and lower bound as well.

The upper bound proposed in equation 2.54 can be seen to hold not only for the theoretical VIX option prices but also for the market prices in all cases except when VIX futures traded above their upper bound. This upper bound as shown in figure 3.1, is given by the maximum of the "F^{VIX} upper bound" curve and "F^{VIX} lower bound" curve. Furthermore, it can be seen that the "F^{VIX} upper bound", "F^{VIX} lower bound" and "F^{VIX} upper market" curves intersect at nearly the same point.

In general, the pricing method over priced in the money options and under priced out of the money options for longer times to maturity (March 14, 21, 28). This can be seen in the figures depicting price and relative errors. If the lognormal assumption was correct, this would imply that the VIX option market expected a higher volatility in the VIX index than that implied by the underlying SPX options. For shorter times to maturity (April 04, 11), the method prices in the money options accurately, but underestimates the prices of out of the money options. Again, if the lognormal assumption was correct, this would imply that the market sees an increase in the VIX as being more likely than than the underlying SPX options imply. This is supported by the fact that VIX futures generally traded close to the upper bound as the time to expiry decreased.

3.5 Concluding remarks

The study can be taken as evidence that the bounds proposed for VIX options hold as long as the VIX futures price is below its theoretical upper bound. If the VIX futures price is above the upper bound, it is possible that it is misprinted relative to the underlying SPX options due to the model free nature of the upper bound. The Carr-Lee method for VIX option pricing produces significant errors relative to market prices for out of the money options. The cause of this discrepancy could be attributed to the failure of the lognormal assumption. Other factors not considered in this paper could also have an effect such as the costs of replicating VIX options via SPX options and/or possible inefficiency in the VIX options market. In order to better identify these sources of discrepancy, further research into both the empirical distribution of future realized volatility and the effects transaction costs would be warranted. Given the popularity of VIX options, it seems that it would be highly desirable to develop a method of replicating VIX options and futures via their underlying SPX options. While the Carr-Lee method attempts to accomplish this, the initial results presented here seem indicate a careful re-evaluation of the underlying assumptions is required.

Appendix A

Pertinent notes

A.1 Results about the quadratic variation of Itô processes

The following applies Øksendal's ([21], Chapter 4) treatment of Itô processes and stochastic integrals to the specific Itô process considered as a model of an asset's price process throughout this paper. The results are then extended to the log process, its quadratic variation and the realized variance of the price process.

Consider a 1-dimensional Wiener process W_t, $t \geq 0$, on the probability space $(\Omega, \mathcal{F}, P)$. We can then define the Itô process S.

Definition 4. An asset price process

$$S_T = S_0 + \int_0^T \mu_t S_t dt + \int_0^T \sigma_t S_t dW_t \tag{A.1}$$

which, in differential form, is

$$dS_t = \mu_t S_t dt + \sigma_t S_t dW_t \tag{A.2}$$

The Itô integral in A.1 is defined if $v(t, \omega) := \sigma_t S_t(\omega)$ satisfies the condition

$$P\left(\int_0^t v(s,\omega)^2 ds < \infty \, for \, all \, t \geq 0\right) = 1 \tag{A.3}$$

and $u(t, \omega) := \mu_t S_t(\omega)$ satisfies the condition

$$P\left(\int_0^t |u(s,\omega)| ds < \infty \, for \, all \, t \geq 0\right) = 1 \tag{A.4}$$

Given the filtration $\mathcal{H}_t \supset \mathcal{F}_t$, with respect to which W_t is a martingale, for our analysis it suffices to consider processes u and v that satisfy the conditions above and are $\mathcal{H}_t$ adapted. Assuming dynamics described by A.1 has two key ramifications: first, the process is continuous, i.e. no jumps, and second, the realized variance of the process must be finite almost surely, that is,

$$\int_0^t \sigma_s^2 ds < \infty \, a.s \, for \, all \, t \geq 0 \tag{A.5}$$

Now we wish to determine the stochastic differential equation for the process $df(t, S_t) := d(\log S_t)$. Applying Itô's formula to A.1, noting that the partial derivatives of f evaluated at (t, s) are $f_t = 0$, $f_s = 1/s$, $f_{ss} = -1/s^2$, we have

$$df(t, S_t) = 0 + \left(\frac{\mu S_t}{S_t} - \frac{\sigma^2 S_t^2}{2 S_t^2} \right) dt + \frac{\sigma S_t}{S_t} dW_t \tag{A.6}$$

which defines the log price process.

Definition 5. The log price process

$$dX_t := d(\log S_t) = \left(\mu_t - \frac{\sigma_t^2}{2} \right) dt + \sigma_t dW_t \tag{A.7}$$

Now we show that the quadratic variation of X_t is the realized variance of A.1. Consider the partition of a time interval $0 = t_0 < t_1 < ... < t_n = T$, and let $M := \max\{|t_i - t_{i-1}| : i = 1, ..., n\}$ be the mesh of that partition such that $||M|| \to 0$ as $n \to \infty$.

Definition 6. The quadratic variation of the process X_t is expressed as the following limit in $L^2(P)$.

$$\langle X \rangle_t = \lim_{n \to \infty} \sum_{i=1}^{n} \left(X_{t_i} - X_{t_{i-1}} \right)^2 \tag{A.8}$$

Proposition 7. *Given A.1 is defined and continuous over the interval $[0, T]$, the quadratic variation of X_t yields the realized variance of S_t.*

$$\langle X \rangle_T = \int_0^T \sigma_t^2 dt \tag{A.9}$$

Proof. The process X_t is continuous and possesses the properties defined for A.1. Define $\langle X \rangle_T$ in terms of its Itô integral

$$\langle X \rangle_t = \lim_{n \to \infty} \sum_{i=1}^{n} \left(X_{t_i} - X_{t_{i-1}} \right)^2 = \int_0^T (dX)^2$$

Using our definition of dX_t

$$(dX_t)^2 = (\cdot)dt \cdot dt + (\cdot)dt \cdot dW_t + \sigma_t^2 dW_t \cdot dW_t$$

From Itô 's formula it is known that $dt \cdot dt = dt \cdot dW = 0$ and $dW \cdot dW = dt$; the shorthand result of the convergence of their respective limit of sums in $L^2(P)$. Thus, $(dX_t)^2 = \sigma_t^2 dt$ and we have the desired result. □

Remark 8. The realized variance of S_t is approximated by

$$\sum_{i=1}^{n} \log^2 \frac{S_{t_i}}{S_{t_{i-1}}} \tag{A.10}$$

for $n \in \mathbb{N}$, which is common practice and gives rise to some approximation error.

A.2 Decomposition of payoffs into European options

The following result is generally attributed to Carr-Madan [4], but several other variants and derivations exist. Proposed here is a simple Taylor series expansion of the payoff function which highlights the model independence of the formula.

Proposition 9. *Decomposition of payoffs into a European option position (Carr-Madan formula)*

Given a contingent T-claim with underlying S and a continuous twice differentiable payoff function $f : \mathbb{R}_+ \mapsto \mathbb{R}$ such that the holder receives $f(S_T)$ at time T, the payoff may be represented as a composition of European call and put options in the following manner

$$f(S_T) = f(\kappa) + f'(\kappa)\,(S_T - \kappa) + \int_0^{\kappa} f''(K)\,(K - S_T)^+\,dK + \int_{\kappa}^{\infty} f''(K)\,(S_T - K)^+\,dK \tag{A.11}$$

where $\kappa \in [0, \infty)$ is arbitrary.

Proof. By a first degree Taylor polynomial of f at κ with an integral remainder

$$f(S_T) = f(\kappa) + f'(\kappa)\,(S_T - \kappa) + \int_{\kappa}^{S_T} f''(K)\,(S_T - K)\,dK$$

The integral remainder can be decomposed such that

$$\int_{\kappa}^{S_T} f''(K)\,(S_T - K)\,dK = \int_{S_T}^{\kappa} f''(K)\,(K - S_T)^+\,dK + \int_{\kappa}^{S_T} f''(K)\,(S_T - K)^+\,dK$$

Since

$$f''(K)\,(K - S_T)^+ = 0 \quad for\, K \leq S_T$$

the lower limit of integration can be changed on the "put" integral

$$\int_{S_T}^{\kappa} f''(K)\,(K - S_T)^+\,dK = \int_{0}^{\kappa} f''(K)\,(K - S_T)^+\,dK$$

Similarly, the upper limit of integration can be changed for the "call" integral

$$\int_{\kappa}^{S_T} f''(K)\,(S_T - K)^+\,dK = \int_{\kappa}^{\infty} f''(K)\,(S_T - K)^+\,dK$$

Combining the two previous results yields the expression in equation A.11. □

Corollary 10. *The risk neutral expectation of A.11 with* $\kappa = F_t$ *is*

$$\mathbb{E}_t\left[f(S_T)\right] = f(F_t) + \int_{0}^{F_t} f''(K)P(K)dK + \int_{F_t}^{\infty} f''(K)C(K)dK$$

where the forward price $F_t = \mathbb{E}_t\left[S_T\right]$*, and* $C(K)$*,* $P(K)$ *are the un-discounted call and put prices of European options expiring at* T*.*

The real power of A.11 is in the the corollary above. Given any continuous, twice differentiable payoff function, the expected value of the associated contingent T-claim can be represented, independent of any model, with market data, assuming European options are traded on the underlying.

A.3 Realized variance and volatility in the Heston model

The Heston model, introduced by Heston in his well known 1993 paper [22], proposes the risk neutral process defined by the following stochastic differential equations:

$$dS_t = rS_t dt + \sqrt{v_t} S_t dW_{1t} \tag{A.12}$$

$$dv_t = \lambda(\bar{v} - v_t) + \eta\sqrt{v_t} dW_{2t} \tag{A.13}$$

where W_1 and W_2 are correlated Wiener processes under the risk neutral measure such that $\mathbb{E}^{\mathbb{Q}}\left[dW_{1t}dW_{2t}\right] = \rho dt$, S is the price process, v is the variance, r is the risk-free rate and λ, $\bar{v}$, η are parameters which can be interpreted as the reversion rate, the mean variance and volatility of volatility. The mean-reverting variance process specified by A.13 was first introduced in finance by Cox, Ingersoll and Ross [23] to model interest rates. They show that the expectation of variance is

$$\mathbb{E}\left[v_t \mid v_0\right] = \bar{v} + (v_0 - \bar{v})e^{-\lambda t}$$

which can be integrated, assuming the integrability of the variance process, to find the expected annualized realized variance:

$$\mathbb{E}\left[R_{0,T}^2\right] := \frac{1}{T}\mathbb{E}\left[\int_0^T v_t dt \mid v_0\right] = \bar{v} + \frac{(v_0 - \bar{v})\left(1 - e^{-\lambda T}\right)}{\lambda T}$$

If $v_0 = \bar{v}$, then the expectation of future variance as well as annualized realized variance is simply $\bar{v}$. The expectation of realized volatility is less straight forward. While the original work is credited to other authors, Gatheral ([19], 144) does a fine job of presenting it in this context. Again we can use the well known formula relating the moment of a random variable to the Laplace transform of its distribution function (as found in Schruger [24]).

$$\mathbb{E}X^r = \frac{1}{\Gamma(1-r)}\int_0^\infty \frac{1 - \mathbb{E}e^{-zX}}{z^{1+r}}dz, \quad 0 < r < 1$$

Applied to relied volatility, i.e. $r = 1/2$, we have:

$$\mathbb{E}R_{0,T} = \frac{1}{2\sqrt{\pi}}\int_0^\infty \frac{1 - \mathbb{E}e^{-zR_{0,T}^2}}{z^{3/2}}dz \tag{A.14}$$

Again, thanks to the work of Cox, Ingersoll and Ross, the Laplace transform we need is already given in [23].

$$\left(\mathcal{L}f_{R^2}\right)(z) = \mathbb{E}e^{-zR_{0,T}^2} = Ae^{-zv_0 B}$$

Where

$$\begin{aligned} A &= \left(\frac{2he^{(\lambda+h)T/2}}{2h + (\lambda+h)(e^{hT}-1)}\right)^{2\lambda\bar{v}/\eta^2} \\ B &= \frac{2(e^{hT}-1)}{2h + (\lambda+h)(e^{hT}-1)} \\ h &= \sqrt{\lambda^2 + 2\eta^2 z} \end{aligned}$$

Equation A.14 can be integrated numerically to provide the expected value of realized volatility under Heston dynamics, but care has to be taken regarding the limits of the computer number format used.

A.4 The VIX call upper bound

The following is given to as support for proposition 3. Using the pricing formula for VIX options given in section 2.4.3, we denote the VIX futures market price as F^{VIX} and a set introduce a variable F which can take any value within the bounds given by equation 2.47, i.e. $\sqrt{A} > F \geq B$. Without loss of generality we set $t = 0$ and suppress $t-$dependence. We also assume $r = 0$. Furthermore we require that the strike, expectation of realized variance, expectation of realized volatility and futures price are positive, that is $K, A, B, F > 0$. The non-negative VIX call price function $C\left(F, K\right)$ (as given in equation 2.53) is required to be monotonically decreasing and convex in K in order for the assumption of no arbitrage to hold as Gatheral [20] points out. Differentiation of $C(F, K)$ with respect to K yields

$$\frac{\partial C}{dK} = \frac{\varphi(d) - \frac{F}{K}\varphi(d+s)}{s} - N(d) \tag{A.15}$$

which can be seen to be less than or equal to zero for $K \in (0, \infty)$. Additionally, $\frac{\partial^2 C}{dK^2} \geq 0$ implying convexity and ensuring that the risk neutral density is non-negative (refer to Gatheral [20]).

Letting $F \to \sqrt{A}$ implies that the variance of the assumed lognormal distribution is zero, so $\sqrt{A} - K = C\left(\sqrt{A}, K\right)$.

Letting $K \to 0$ we see that $\sqrt{A} > C\left(F^{VIX}, 0\right) \geq C\left(B, 0\right)$. Now we wish to show that proposition 3 holds, i.e.

$$C\left(F^{VIX}, K\right) \leq \max\left\{\sqrt{A_t} - K, C\left(B_t, K\right)\right\} \tag{A.16}$$

Consider the partial derivative the call price (equation 2.53) with respect to the VIX futures price:

$$\frac{\partial C}{dF} = N(d+s) + \frac{d+s}{s^2}\varphi(d+s) - \frac{K\varphi(d)}{Fs^2}\left(2s+d\right) \tag{A.17}$$

where

$$\begin{aligned} d &= \frac{\mu - \log K}{s} & \text{(A.18)} \\ \mu &= 2\log F - \frac{1}{2}\log A & \text{(A.19)} \\ s^2 &= \log A - 2\log F & \text{(A.20)} \end{aligned}$$

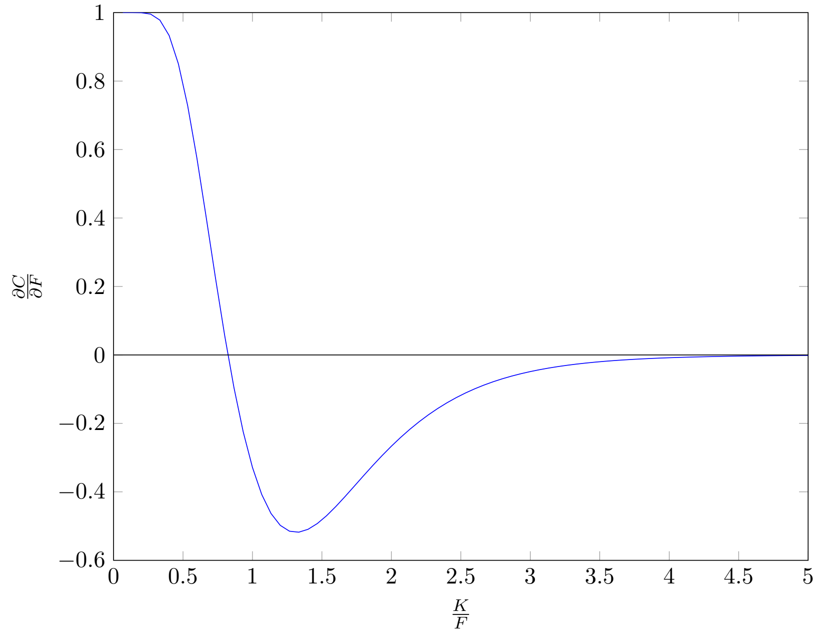

FIGURE A.1: The sensitivity of call price to futures value, $\frac{\partial C}{\partial F}$, for $\sqrt{A} = 1.1F$

as before. The partial derivative $\frac{\partial C}{\partial F}$ is plotted as a function of K/F in figure A.1, setting the upper bound $\sqrt{A} = cF$ for constant F. It can be seen that for any F within the prescribed bounds (i.e. $c > 1$ so that $\sqrt{A} > F$), $\frac{\partial C}{\partial F} > 0$ for $K < K_\star$ and $\frac{\partial C}{\partial F} < 0$ for $K > K_\star$, where $K_\star$ is the strike satisfying

$$\frac{\partial C}{\partial F}(F, K_\star) = 0 \tag{A.21}$$

Defining $K_\star$ as the strike where

$$\frac{\partial C}{\partial F}\left(F^{VIX}, K_\star\right) = 0 \tag{A.22}$$

and given the bounds on F^{VIX}, we have $C(B, K) \geq C\left(F^{VIX}, K\right)$ for $K \geq K_\star$and $\sqrt{A} - K \geq C\left(F^{VIX}, K\right)$ for $K \leq K_\star$, meaning that $C\left(F^{VIX}, K\right)$ is dominated by either $C(B, K)$ or $\sqrt{A} - K$ for all K, which yields the upper bound given in equation 2.54.

Appendix B

Additional plots

The curve marked “F^{VIX} market” was generated by applying the formula proposed by Carr-Lee, given in equation 2.53. The curve marked “F^{VIX} upper bound” was generated by replacing the market price F^{VIX} with its upper bound in the valuation formula 2.53. Similarly, the curve marked “F^{VIX} lower bound” was generated by replacing F^{VIX} with its lower bound and the curve marked “mean” was generated by replacing F^{VIX} with the mean of the lower bound and upper bound.

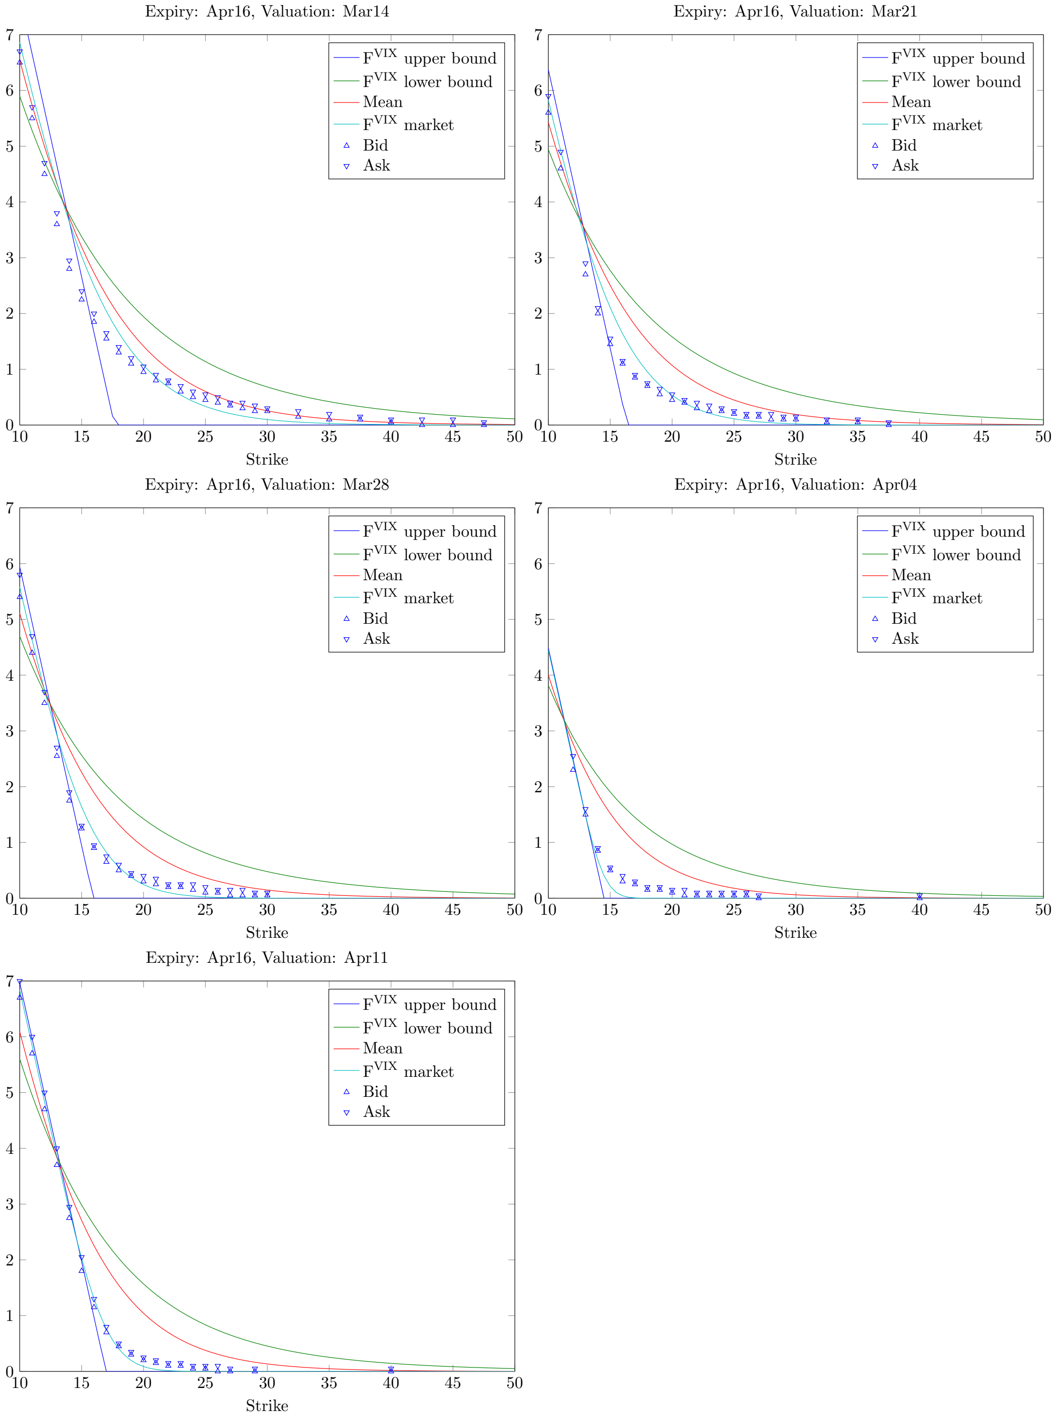

FIGURE B.1: April 16 expiry: comparison of theoretical and market VIX options prices

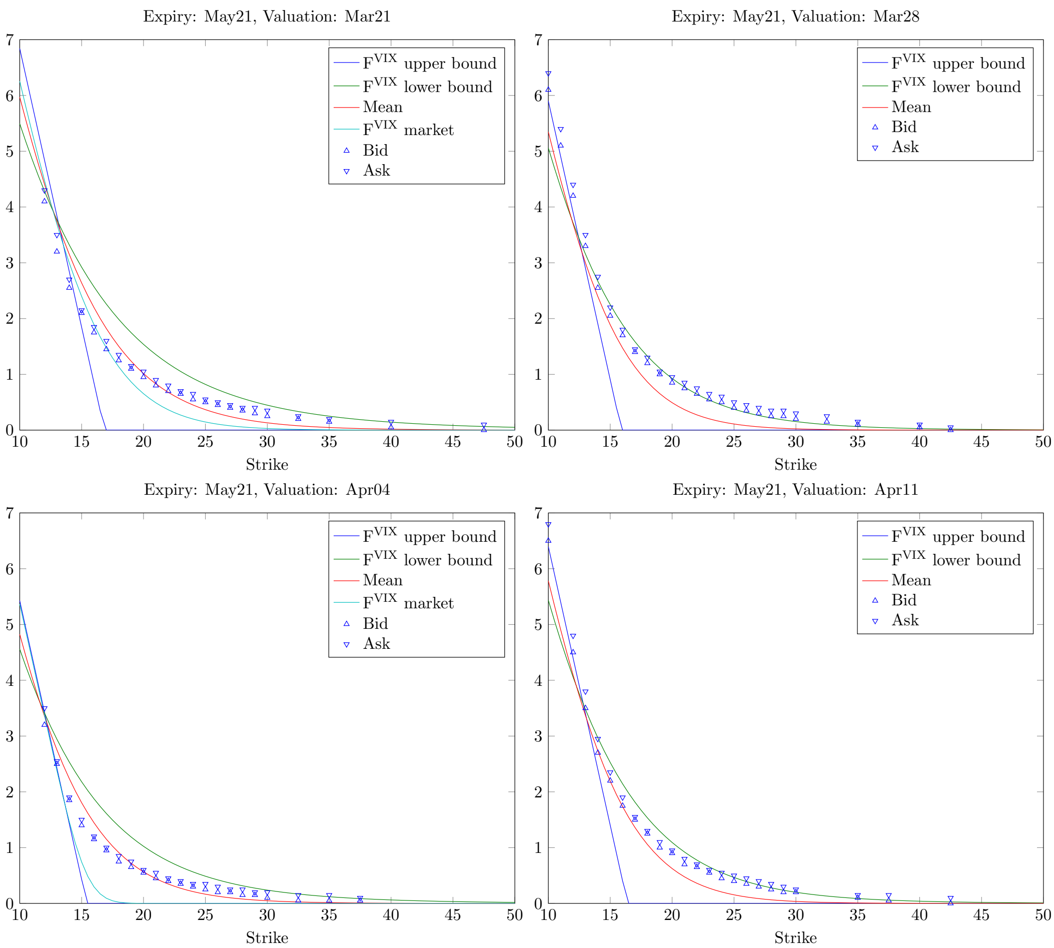

FIGURE B.2: May 21 expiry: comparison of theoretical and market VIX options prices

Bibliography

[1] Harry Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.

[2] Anthony J Neuberger. *Volatility trading.* Institute of Finance and Accounting, London Business School, 1990.

[3] Bruno Dupire. Model art. *Risk*, 6(9):118–124, 1993.

[4] Peter Carr and Dilip Madan. Towards a theory of volatility trading. *Option Pricing, Interest Rates and Risk Management, Handbooks in Mathematical Finance*, pages 458–476, 2001.

[5] Peter Carr and Roger Lee. Volatility derivatives. *Annu. Rev. Financ. Econ.*, 1(1): 319–339, 2009.

[6] Chicago Board Options Exchange. The cboe volatility index–VIX. *White Paper*, 2009.

[7] Lewis Biscamp and Tim Weithers. Variance swaps and CBOE S&P 500 variance futures. *Unpublished working paper, Chicago Trading Company, LLC*, 2007.

[8] Zhiguang Wang and Robert T Daigler. The performance of vix option pricing models: empirical evidence beyond simulation. *Journal of Futures Markets*, 31(3): 251–281, 2011.

[9] Robert E Whaley. Derivatives on market volatility: Hedging tools long overdue. *The journal of Derivatives*, 1(1):71–84, 1993.

[10] Andreas Grünbichler and Francis A Longstaff. Valuing futures and options on volatility. *Journal of Banking & Finance*, 20(6):985–1001, 1996.

[11] Peter Carr and Roger Lee. Robust replication of volatility derivatives. In *PRMIA award for Best Paper in Derivatives, MFA 2008 Annual Meeting*, 2008.

[12] Kresimir Demeterfi, Emanuel Derman, Michael Kamal, and Joseph Zou. More than you ever wanted to know about volatility swaps. *Goldman Sachs Quantitative Strategies Research Notes*, 1999.

[13] Peter Carr and Roger Lee. Realized volatility and variance: Options via swaps. *Risk*, 20(5):76–83, 2007.

[14] Steven Phillip Feinstein. *The Black-Scholes formula is nearly linear in sigma for at-the-money options: Therefore implied volatilities from at-the-money options are virtually unbiased*. Federal Reserve Bank of Atlanta, 1989.

[15] Menachem Brenner and Marti G Subrahmanyam. A simple formula to compute the implied standard deviation. *Financial Analysts Journal*, pages 80–83, 1988.

[16] Peter Friz and Jim Gatheral. Valuation of volatility derivatives as an inverse problem. *Quantitative Finance*, 5(6):531–542, 2005.

[17] Gurdip Bakshi, Charles Cao, and Zhiwu Chen. Empirical performance of alternative option pricing models. *The Journal of Finance*, 52(5):2003–2049, 1997.

[18] Fabrice D Rouah. *The Heston Model and Its Extensions in Matlab and C#*. John Wiley & Sons, 2013.

[19] Jim Gatheral. *The volatility surface: a practitioner's guide*, volume 357. John Wiley & Sons, 2006.

[20] Jim Gatheral and Antoine Jacquier. Arbitrage-free svi volatility surfaces. *Quantitative Finance*, 14(1):59–71, 2014.

[21] Bernt Øksendal. *Stochastic differential equations*. Springer, 2003.

[22] Steven L Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of financial studies*, 6(2):327–343, 1993.

[23] John C Cox, Jonathan E Ingersoll Jr, and Stephen A Ross. A theory of the term structure of interest rates. *Econometrica: Journal of the Econometric Society*, pages 385–407, 1985.

[24] Klaus Schürger. Laplace transforms and suprema of stochastic processes. In *Advances in Finance and Stochastics*, pages 285–294. Springer, 2002.