Application of Volatility Targeting Strategies within a Black-Scholes Framework

Dmitri Vakaloudis

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in partial fulfilment of the requirements for the degree of Master of Philosophy.

October 13, 2019

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Philosophy in the University of Cape Town. It has not been submitted before for any degree or examination to any other University.

Signed by candidate

October 13, 2019

Abstract

The traditional Black-Scholes (BS) model relies heavily on the assumption that underlying returns are normally distributed. In reality however there is a large amount of evidence to suggest that this assumption is weak and that actual return distributions are non-Gaussian. This dissertation looks at algorithmically generating a Volatility Targeting Strategy (VTS) which can be used as an underlying asset. The rationale here is that since the VTS has a constant prespecified level of volatility, its returns should be normally distributed, thus tending closer to an underlying that adheres to the assumptions of BS.

Acknowledgements

Firstly I would like to thank the African Institute of Financial Markets and Risk Management (AIFMRM) for giving me the opportunity to participate in such an enriching course. Secondly a special word of thanks needs to go out to my supervisor Obeid Mahomed for all the hours of valuable insight he offered over the course of this dissertation. Finally I would like to thank Greg Mollentz from Old Mutual Specialised Finance for providing such an interesting topic.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Research Overview

Liquid options are quoted in volatility rather than monetary terms with the understanding that the price of the derivative can be obtained by inserting this volatility into the BS formulae. This measure of volatility is known as implied volatility and when plotted against strike and maturity, non-constant features such as skews or smiles are often observed. A smile reflects fat tails in the return distribution whereas a skew indicates return distribution asymmetry. Further testament to this asymmetry is the inflative nature of volatility over time, known as the term structure of implied volatility. Such features are explicitly defined by Homescu (2011).

The existence of skew, smile and term structure introduces several complications of which the BS model cannot assimilate. Firstly pricing and hedging of options becomes more complicated as volatility is obviously no longer assumed to be constant. These inaccuracies become more prevalent as option complexity increases. Secondly arbitrage may exist among options not quoted in the market (where prices of these unquoted options are obtained from interpolating or extrapolating the implied volatility surface) even if the quoted market set is arbitrage free. As a result put-call parity no longer holds implying volatility quotes are no longer applicable for both calls and puts further complicating the process of pricing and risk managing options.

A VTS aims at algorithmically generating an asset with a constant prespecified volatility level. The rationale behind such a strategy is that since this underlying has a constant volatility, its returns should be normally distributed, thus tending closer to a BS world. In a practical sense, the VTS is an active portfolio allocation strategy that generates a risk-adjusted underlying.

In conventional active portfolio management, components are selected for the purpose of obtaining alpha while minimizing beta, or earning returns over a market benchmark while minimizing systematic market volatility. This is achieved

through asset diversification across classes and/or geographical locations. During periods of market correction, assets selected for their uncorrelated behaviour suddenly become correlated as market fear intensifies. Investments selected for their ability to generate alpha without beta begin to exhibit the opposite behaviour. Hocquard *et al.* (2013) stated that this phase-locking behaviour, coupled with jumps in market volatility result in large draw-downs which limit expected returns of these actively managed portfolios.

The VTS (X_t) eliminates the manual process of allocation by targeting volatility (σ_{target}) instead of returns. This is achieved through the process of making portfolio weightings (ω_t) dependent on the volatility inherent in the risky component of the portfolio. This weighting, known as the participation ratio, can be expressed as the desired level of volatility over the actual measured level. Therefore the fundamental notion behind the VTS is that the volatility of the risky component of the portfolio can be measured and managed.

If an accurate estimation of this volatility can be obtained, then a hedge can be implemented (in the form of the money market trade) that can adjust the total volatility inherent in the portfolio. This is of course based on the assumption that the money market account is analogous to a risk-free asset. The hedging strategy works by dynamically managing holdings in a risky and risk-free asset. There are three possibilities for hedging:

1. If the desired level of volatility is less than the actual measured level, sell the risky asset and use the proceeds to invest in the money market account

2. If the desired level of volatility is greater than the actual measured level, buy the risky asset by selling a portion of the money market account

3. If the desired level of volatility is equal to the actual measured level, the investor's risk appetite is equal to the risk observed in the market, therefore only invest in the risky asset

By continuously conducting this hedge over the life of the investment strategy, a portfolio with constant volatility characteristics should be obtained, i.e. a portfolio that follows Geometric Brownian Motion (GBM). Of course in reality it is neither possible or financially feasible to continuously hedge over time and so discrete hedging would need to be implemented, introducing drag into the portfolio. Drag is defined as the practical costs incurred from rebalancing the portfolio, such as transactional trading costs and broker fees. If portfolio drag is high, profits could be compromised regardless of the performance of the portfolio's assets.

The dynamics of the VTS can be expressed mathematically as

$$\frac{\mathrm{d}X_t}{X_t} = \omega_t \left(\frac{\mathrm{d}S_t}{S_t}\right) + (1-\omega_t)\left(\frac{\mathrm{d}B_t}{B_t}\right) \tag{1.1}$$

in continuous time or

$$\frac{X_{t+\Delta t} - X_t}{X_t} = \omega_t \left(\frac{S_{t+\Delta t} - S_t}{S_t}\right) + (1-\omega_t)\left(\frac{B_{t+\Delta t} - B_t}{B_t}\right) \tag{1.2}$$

in discrete time (in terms of percentage returns) over time interval Δt. S_t denotes the risky asset whereas B_t denotes the risk-free asset or money market account. S_t could comprise of many different assets such as an Exchange-Traded Fund (ETF).

Papageorgiou *et al.* (2017) implemented a VTS where a Generalized Autoregressive Conditional Heteroscedasticity (GARCH) model was used to estimate the actual volatility inherent in the risky asset. This dissertation builds on this by using stochastic volatility estimates instead of statistical.

Even though at times volatility is not constant, consistently predictable or even directly observable, several stylised facts have emerged over the years with regard to its behaviour which make it a prime candidate for stochastic modelling.

Empirical studies have shown that for most asset classes volatility, no matter how extreme at any point in time, gradually returns to a mean level through a process known as mean reversion. Often periods of extreme volatility persist and it takes a while for the underlying to recover to a normal level, giving rise to the notion of volatility persistence or clustering. In other words periods where volatility is either high or low are preceded by periods where volatility remains high or low. These characteristics were first observed by Mandelbrot (1997) and Fama (1998).

Implied volatility serves as the market's perception of how risky an option's underlying may be. When stock markets crash, option implied volatilities tend to increase and vice versa. This is intuitive because as stock prices decrease, fear that the price may decrease further translates into higher quoted volatility. This phenomenon is especially prevalent for deep out-the-money (OTM) options which are usually priced cheaply and is directly responsible for volatility skew or smile features already mentioned.

The final stylised fact is the presence of volatility jumps. Jumps represent sharp movements in the underlying that occur over a very short period of time. Socio-economic events or announcements are the main trigger for these jumps and in general there are a higher frequency of downward directed jumps than upward, originating from the risk averse nature of most investors. This behaviour was first observed by Bates (1996).

Stochastic Volatility Models (SVMs) aim at treating volatility as a stochastic variable itself. Over the years various SVMs have been developed with the goal

of capturing a wide range of market dynamics. This dissertation focuses on two models, both capable of assimilating the aforementioned stylised facts, namely the Heston and Bates models.

1.2 Research Method & Aims

The primary goal of this dissertation is to construct a VTS portfolio based on stochastic volatility estimates. Here the VTS is constructed for the purpose of investing and is derived under the real-world measure.

The secondary goal is to use this constructed portfolio as an underlying that options can be written on. Here the VTS is derived under the risk-neutral measure. This application is useful from an insurance perspective where exotic, path-dependent options are issued to clients forming part of a pension or retirement annuity product. They are popular for the level of guarantee they can offer over long investment horizons. Here it is important that the underlying's distributional and volatility characteristics are well behaved for the entirety of the contract. These behavioural characteristics are key to the success of the VTS if it were to be implemented practically.

The following assumptions were made in this dissertation:

1. A risk-free rate exists and remains constant over time, implying there is no correlation between the risky and risk-free asset
2. The risky asset pays zero dividends
3. Costs incurred to rebalance the portfolio are considered negligible as the rate is kept constant (at daily rebalancing) for all simulations
4. The risky asset can be perfectly modelled using the the Heston and Bates models, implying no model risk. This assumption is two fold;
    - 4.1. Firstly it is assumed that the risky asset behaves in accordance with the volatility stylized facts previously mentioned. Therefore this dissertation does not limit the type of asset that can be used as the risky asset, so long as it pays no dividends and exhibits volatility characteristics inline with the stylized facts
    - 4.2. Secondly the Heston and Bates models perfectly model the aforementioned stylized facts and their model parameters can be obtained through calibration. Implying there is a complete set of option prices, written on the chosen risky asset, that can recover the measured level of volatility observed in the market

The proceeding chapters are structured as follows:

Chapter 2 defines the mathematical derivation of the VTS, using both the Heston and Bates SVMs to model the risky asset. From this point onwards the time-indexed VTS driven by the aforementioned models will be referred to as X_t^H and X_t^B respectively. Initially the strategies are defined in continuous time, where their distributional properties are studied. Thereafter the Euler discretization scheme is used to discretize the risky and risk-free asset, as well as the variance dynamics of the risky asset. Here the notion of a hedging error is introduced as the error brought about by how frequently the participation ratio is adjusted through calibration. Next the risk-neutral dynamics of the risky asset are defined for the Heston and Bates models as well as the variance process.

Chapter 3 defines how the VTS could be constructed practically for both investment and option writing purposes. The hedging error is tested for different calibration rates to test how a practically constructed VTS would differ from the continuous solution. Chapter 3 is concluded with a MCS where vanilla options, written on the VTS, are priced in terms of implied volatility to illustrate the performance of the VTS as an underlying for several strikes and maturities. For at-the-money (ATM) vanilla options, several calibration rates are tested to illustrate the impact of discretization on option pricing.

Chapter 4 uses the VTS as an underlying on which an exotic path-dependent lookback put with floating strike is priced. The price is obtained from an error corrected MCS and compared to a closed-form price that assumes returns are normally distributed. Daily and weekly calibration rates are tested here along with various target volatility levels.

Chapter 5 concludes the dissertation and provides areas of future research for the improvement of the VTS.

Chapter 2

Mathematical Specification of the VTS

2.1 Construction under the Real-World Measure

To construct the continuous form of the VTS, under the real-world measure, it is assumed that there exists a time-series of calibrated model parameters that perfectly fit the market at all instantaneous points in time. Under this assumption the Heston and Bates model parameters can be used to simulate the risky asset's price dynamics.

2.1.1 VTS modelled by Heston Model Dynamics

Heston (1993) proposed a SVM where the risky asset follows a BS-type stochastic process. Unlike the conventional BS model, which assumes that the asset's volatility remains constant over time, the Heston model assumes that volatility is driven by a stochastic variance (V_t) which follows a Cox-Ingersoll-Ross (CIR) process.

The real-world dynamics of the model are given by the bivariate set of Stochastic Differential Equations (SDEs)

$$\frac{\mathrm{d}S_t}{S_t} = \mu_t \mathrm{d}t + \sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{1,t} \tag{2.1}$$

$$\mu_t = r + \lambda_s \sqrt{V_t}$$

$$\mathrm{d}V_t = \kappa(\bar{V} - V_t)\mathrm{d}t + \sigma\sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{2,t} \tag{2.2}$$

$$\mathbb{C}\mathrm{ov}\left[\mathrm{d}W^{\mathbb{P}}_{1,t}, \mathrm{d}W^{\mathbb{P}}_{2,t}\right] = \rho \mathrm{d}t \tag{2.3}$$

The drift of the asset (μ_t) can be expressed as the risk-free rate (r) plus the excess rate of return, which is simply the market price of risk (λ_s) multiplied by the time evolving volatility of the asset $\left(\sqrt{V_t}\right)$.

The rationale behind using a CIR process to model volatility is that it is mean reverting. Parameter κ controls the rate at which the process returns to its long run average $\left(\bar{V}\right)$. Kienitz and Wetterau (2012) showed empirically that high values of κ and $\bar{V}$ have an inverse effect on the kurtosis of the return distributions.

$W_{1,t}^{\mathbb{P}}$ and $W_{2,t}^{\mathbb{P}}$ are two standard Brownian motion processes under the real-world measure $\mathbb{P}$ with correlation ρ. Correlation is directly responsible for affecting the skew of the return distribution, with positive values of ρ skewing the distribution to the left and negative values skewing the distribution to the right.

The volatility of variance (σ) is directly responsible for affecting the kurtosis of the distribution. If $2\kappa\bar{V} \geqslant \sigma^2$ then the Feller condition holds and the variance process will never drop below 0 in continuous time.

To simulate the risk-free asset's dynamics, it is assumed that the money market account accrues at the NACC risk-free rate

$$\frac{\mathrm{d}B_t}{B_t} = r\mathrm{d}t \tag{2.4}$$

Substituting equations (2.1) and (2.4) into equation (1.1) while noting that the participation ratio is equal to

$$\omega_t = \frac{\sigma_{target}}{\sqrt{V_t}}$$

since$\sqrt{V_t}$ is the only source of volatility inherent in equation (2.1), yields

$$\begin{aligned}\frac{\mathrm{d}X_t}{X_t} &= \omega_t\left[\left(r+\lambda_s\sqrt{V_t}\right)\mathrm{d}t + \sqrt{V_t}\mathrm{d}W_{1,t}^{\mathbb{P}}\right] + (1-\omega_t)\left(r\mathrm{d}t\right)\\ &= \left(r+\lambda_s\sigma_{target}\right)\mathrm{d}t + \sigma_{target}\mathrm{d}W_{1,t}^{\mathbb{P}}\end{aligned} \tag{2.5}$$

A closed-form solution for equation (2.5) can be obtained by setting $f(t,X_t) := \ln X_t$. Using Itō's Lemme over time interval $t \leq u \leq t+\Delta t$ while noting that, $\frac{\partial f(t,X_t)}{\partial t} := 0$, $\frac{\partial f(t,X_t)}{\partial X_t} := \frac{1}{X_t}$ and $\frac{\partial^2 f(t,X_t)}{\partial X_t^2} := -\frac{1}{X_t^2}$, yields

$$\begin{aligned}\mathrm{d}f(t,X_t) &= \frac{\partial f\left(t,X_t\right)}{\partial t}\mathrm{d}t + \frac{\partial f\left(t,X_t\right)}{\partial X_t}\mathrm{d}X_t + \frac{1}{2}\frac{\partial^2 f\left(t,X_t\right)}{\partial X_t^2}(\mathrm{d}X_t)^2\\ &= \frac{1}{X_t}\mathrm{d}X_t - \frac{1}{2}\left(\frac{1}{X_t^2}\right)(\mathrm{d}X_t)^2\\ &= \frac{1}{X_t}\left[\left(r+\lambda_s\sigma_{target}\right)X_t\mathrm{d}t + \sigma_{target}X_t\mathrm{d}W_{1,t}^{\mathbb{P}}\right] - \frac{1}{2X_t^2}\sigma_{target}^2X_t^2\mathrm{d}t\\ &= \left(r+\sigma_{target}\left(\lambda_s-\frac{1}{2}\sigma_{target}\right)\right)\mathrm{d}t + \sigma_{target}\mathrm{d}W_{1,t}^{\mathbb{P}}\\ f\left(t+\Delta t, X_{t+\Delta t}\right) &= f\left(t,X_t\right) + \int_t^{t+\Delta t}\left(r+\sigma_{target}\left(\lambda_s-\frac{1}{2}\sigma_{target}\right)\right)\mathrm{d}u\\ &\quad + \int_t^{t+\Delta t}\sigma_{target}\mathrm{d}W_{1,u}^{\mathbb{P}}\\ \ln X_{t+\Delta t} &= \ln X_t + \left(r+\sigma_{target}\left(\lambda_s-\frac{1}{2}\sigma_{target}\right)\right)\Delta t + \sigma_{target}W_{1,\Delta t}^{\mathbb{P}}\\ X_{t+\Delta t} &= X_t\exp\left[\left(r+\sigma_{target}\left(\lambda_s-\frac{1}{2}\sigma_{target}\right)\right)\Delta t + \sigma_{target}\sqrt{\Delta t}Z_1\right]\end{aligned} \tag{2.6}$$

where $W_{1,\Delta t}^{\mathbb{P}} = \sqrt{\Delta t}Z_1$ and Z_1 is a standard normal random variate.

To determine the distributional properties of equation (2.5), expectations and variances over the filtered probability space $\left(\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F})_{t\geq 0}\right)$ are calculated. Taking the expectation of both sides with respect to the filtration $\mathcal{F}_t$ yields

$$\begin{aligned}
\mathbb{E}^{\mathbb{P}}\left[\frac{\mathrm{d}X_t}{X_t}\middle|\mathcal{F}_t\right] &= \mathbb{E}^{\mathbb{P}}\left[(r+\lambda_s\sigma_{target})\,\mathrm{d}t + \sigma_{target}\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] \\
&= \mathbb{E}^{\mathbb{P}}\left[(r+\lambda_s\sigma_{target})\,\mathrm{d}t|\mathcal{F}_t\right] + \mathbb{E}^{\mathbb{P}}\left[\sigma_{target}\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] \\
&= (r+\lambda_s\sigma_{target})\,\mathrm{d}t + \sigma_{target}\mathbb{E}^{\mathbb{P}}\left[\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] \\
&= (r+\lambda_s\sigma_{target})\,\mathrm{d}t
\end{aligned}$$

since $\mathrm{d}W^{\mathbb{P}}_{1,t} \sim \mathcal{N}(0, \mathrm{d}t)$. Similarly taking the variance of both sides yields

$$\begin{aligned}
\mathbb{V}\mathrm{ar}\left[\frac{\mathrm{d}X_t}{X_t}\middle|\mathcal{F}_t\right] &= \mathbb{V}\mathrm{ar}\left[(r+\lambda_s\sigma_{target})\,\mathrm{d}t + \sigma_{target}\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] \\
&= \mathbb{E}^{\mathbb{P}}\left[(\sigma_{target}\mathrm{d}W^{\mathbb{P}}_{1,t})^2\middle|\mathcal{F}_t\right] - \left(\mathbb{E}^{\mathbb{P}}\left[\sigma_{target}\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right]\right)^2 \\
&= \mathbb{E}^{\mathbb{P}}\left[\sigma^2_{target}\mathrm{d}t\middle|\mathcal{F}_t\right] - \sigma^2_{target}\left(\mathbb{E}^{\mathbb{P}}\left[\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right]\right)^2 \\
&= \sigma^2_{target}\mathrm{d}t
\end{aligned}$$

Hence one can deduce that the continuous returns of X^H_t have the following distribution under the real-world measure $\mathbb{P}$

$$\frac{\mathrm{d}X_t}{X_t} \sim \mathcal{N}\left((r+\lambda_s\sigma_{target})\,\mathrm{d}t, \sigma^2_{target}\mathrm{d}t\right) \tag{2.7}$$

2.1.2 VTS modelled by Bates Model Dynamics

Bates (1996) proposed an extension of the Heston model which includes jumps in the price process of the risky asset. Equations (2.2) and (2.3) remain applicable for the Bates model, however the real-world risky asset dynamics are given by the following SDE

$$\frac{\mathrm{d}S_t}{S_t} = (\mu_t - \alpha)\,\mathrm{d}t + \sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{1,t} + (J-1)\mathrm{d}N_t \tag{2.8}$$

$$\mu_t = r + \lambda_s\sqrt{V_t + \beta}$$

where α is the drift and β is the variance introduced from the jump process.

$$\begin{aligned}
\alpha\mathrm{d}t &= \mathbb{E}^{\mathbb{P}}\left[(J-1)\mathrm{d}N_t|\mathcal{F}_t\right] \\
&= \mathbb{E}^{\mathbb{P}}\left[(J-1)|\mathcal{F}_t\right]\mathbb{E}^{\mathbb{P}}\left[\mathrm{d}N_t|\mathcal{F}_t\right] \\
&= \left(\exp\left[\mu_J + \frac{1}{2}\sigma^2_J\right] - 1\right)\theta\mathrm{d}t \\
\alpha &= \left(\exp\left[\mu_J + \frac{1}{2}\sigma^2_J\right] - 1\right)\theta
\end{aligned}$$

$$\begin{aligned}\beta \mathrm{d}t &= \mathbb{V}\mathrm{ar}\left[(J-1)\mathrm{d}N_t|\mathcal{F}_t\right] \\ &= \mathbb{V}\mathrm{ar}\left[(J-1)|\mathcal{F}_t\right]\mathbb{V}\mathrm{ar}\left[\mathrm{d}N_t|\mathcal{F}_t\right] \\ &= \left[\left(\exp\left[\sigma_J^2\right]-1\right)\exp\left[2\mu_J+\sigma_J^2\right]\right]\theta\mathrm{d}t \\ \beta &= \left[\left(\exp\left[\sigma_J^2\right]-1\right)\exp\left[2\mu_J+\sigma_J^2\right]\right]\theta\end{aligned}$$

Synonymous with the Heston model, the drift can again be expressed as the risk-free rate plus the excess rate of return, where now the volatility of the risky asset is equal to$\sqrt{V_t+\beta}$ to include the volatility introduced from jumps.

$J = \exp(Z_3)$ represents the jump size magnitude under the log-normal assumption of jumps with Z_3 being a normal random variate with distribution $Z_3 \sim \mathcal{N}\left(\mu_J, \sigma_J^2\right)$. It is interesting to note that $J-1 \in [-100\%, \infty)$ is the percent jump size. For example if Z_3 = -0.3567 = ln(0.7) < 0 then $\exp(Z_3)$ = 0.7 and $\exp(Z_3)$ - 1 = -0.3 therefore the stock jumps down by 30%. It can be deduced that if $Z_3 < 0$ then the stock jumps down and conversely if $Z_3 > 0$ the stock jumps up.

From the above result it is easy to see how μ_J directly effects the skewness of the return distribution. Negative values will result in more jumps down thus skewing the distribution to the left with the opposite being true for positive values.

The number of jumps in each period are defined by a homogeneous Poisson process (N_t). The jump intensity parameter (θ) and volatility of jumps (σ_J) inversely effect the kurtosis of the distribution with higher values resulting in a lower kurtosis. Parameters ρ and σ have the same effect on the return distribution as the Heston model.

Substituting equations (2.8) and (2.4) into equation (1.1) while noting that the participation ratio applicable for the Bates model is equal to

$$\omega_t = \frac{\sigma_{target}}{\sqrt{V_t+\beta}}$$

yields

$$\begin{aligned}\frac{\mathrm{d}X_t}{X_t} &= \omega_t\left[\left(r+\lambda_s\sqrt{V_t+\beta}-\alpha\right)\mathrm{d}t + \sqrt{V_t}\mathrm{d}W_{1,t}^{\mathbb{P}} + (J-1)\mathrm{d}N_t\right] + (1-\omega_t)(r\mathrm{d}t) \\ &= (r+\lambda_s\sigma_{target})\,\mathrm{d}t + \omega_t\sqrt{V_t}\mathrm{d}W_{1,t}^{\mathbb{P}} + \omega_t\left[(J-1)\mathrm{d}N_t - \alpha\mathrm{d}t\right] \qquad (2.9)\end{aligned}$$

The returns defined by equation (2.9) are no longer normally distributed making the process of finding a closed-form solution for X_t^B non-trivial. For this reason the numerical approximation defined in section 2.2 serves as an approximate solution.

If β were ignored in the derivation of equation (2.9), the dynamics would reduce to the Merton model dynamics where the Brownian motion driven volatility of the VTS would remain constant at σ_{target}. This might seem like a desirable result but is overlooked due to the fact that the volatility introduced from the Poisson process would be ignored and therefore not compensated for.

To determine the distributional properties of equation (2.9), expectations and variances over the filtered probability space $\left(\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F})_{t\geq 0}\right)$ are again calculated.

Taking the expectation of both sides with respect to the filtration $\mathcal{F}_t$ yields

$$\begin{aligned}\mathbb{E}^{\mathbb{P}}\left[\frac{\mathrm{d}X_t}{X_t}\middle|\mathcal{F}_t\right] &= \mathbb{E}^{\mathbb{P}}\left[(r+\lambda_s\sigma_{target})\,\mathrm{d}t + \omega_t\sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{1,t} + \omega_t\left[(J-1)\mathrm{d}N_t - \alpha\mathrm{d}t|\mathcal{F}_t\right]\right] \\ &= \mathbb{E}^{\mathbb{P}}\left[(r+\lambda_s\sigma_{target})\,\mathrm{d}t|\mathcal{F}_t\right] + \mathbb{E}^{\mathbb{P}}\left[\omega_t\sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] \\ &+ \mathbb{E}^{\mathbb{P}}\left[\omega_t\left[(J-1)\mathrm{d}N_t - \alpha\mathrm{d}t\right]|\mathcal{F}_t\right] \end{aligned} \tag{2.10}$$

where

$$\begin{aligned}\mathbb{E}^{\mathbb{P}}\left[\omega_t\left[(J-1)\mathrm{d}N_t - \alpha_t\mathrm{d}t\right]|\mathcal{F}_t\right] &= \mathbb{E}^{\mathbb{P}}\left[\omega_t|\mathcal{F}_t\right]\mathbb{E}^{\mathbb{P}}\left[(J-1)\mathrm{d}N_t - \alpha\mathrm{d}t|\mathcal{F}_t\right] \\ &= \omega_t\left(\alpha\mathrm{d}t - \mathbb{E}^{\mathbb{P}}\left[\alpha\mathrm{d}t|\mathcal{F}_t\right]\right) \\ &= 0\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}^{\mathbb{P}}\left[\omega_t\sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] &= \omega_t\sqrt{V_t}\mathbb{E}^{\mathbb{P}}\left[\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] \\ &= 0\end{aligned}$$

since ω_t and$\sqrt{V_t}$ are $\mathcal{F}_t$ measurable and therefore constants. Equation (2.10) simply becomes

$$\mathbb{E}^{\mathbb{P}}\left[\frac{\mathrm{d}X_t}{X_t}\middle|\mathcal{F}_t\right] = (r+\lambda_s\sigma_{target})\,\mathrm{d}t$$

Similarly taking the variances of both sides of equation (2.9) yields

$$\begin{aligned}\mathbb{V}\mathrm{ar}\left[\frac{\mathrm{d}X_t}{X_t}\middle|\mathcal{F}_t\right] &= \mathbb{V}\mathrm{ar}\left[(r+\lambda_s\sigma_{target})\,\mathrm{d}t + \omega_t\sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{1,t} + \omega_t\left[(J-1)\mathrm{d}N_t - \alpha\mathrm{d}t|\mathcal{F}_t\right]\right] \\ &= \mathbb{V}\mathrm{ar}\left[(r+\lambda_s\sigma_{target})\,\mathrm{d}t|\mathcal{F}_t\right] + \mathbb{V}\mathrm{ar}\left[\omega_t\sqrt{V_t}\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] \\ &+ \mathbb{V}\mathrm{ar}\left[\omega_t\left[(J-1)\mathrm{d}N_t - \alpha\mathrm{d}t\right]|\mathcal{F}_t\right] \\ &= \omega_t^2 V_t\mathbb{V}\mathrm{ar}\left[\mathrm{d}W^{\mathbb{P}}_{1,t}\middle|\mathcal{F}_t\right] + \omega_t^2\mathbb{V}\mathrm{ar}\left[(J-1)\mathrm{d}N_t|\mathcal{F}_t\right] \\ &= \omega_t^2 V_t\mathrm{d}t + \omega_t^2\beta\mathrm{d}t \\ &= \omega_t^2\left(V_t+\beta\right)\mathrm{d}t \\ &= \sigma^2_{target}\mathrm{d}t\end{aligned}$$

Therefore the continuous returns of X_t^B have the following distribution

$$\frac{\mathrm{d}X_t}{X_t} \sim \mathcal{D}\left((r+\lambda_s\sigma_{target})\,\mathrm{d}t, \sigma^2_{target}\mathrm{d}t\right) \tag{2.11}$$

which is equivalent to X_t^H under the real-world measure $\mathbb{P}$ in terms of mean and variance magnitudes, but the underlying return distribution $\mathcal{D}$ is unknown.

2.2 Discretization & Incurred Hedging Errors

The above derivation assumes that the portfolio can be hedged continuously, i.e. at any instantaneous point in time calibration occurs, the participation ratio is calculated and an appropriate hedge can be implemented to compensate for the time-varying volatility. In reality however the participation ratio can only be calculated over discrete time intervals. Over time interval δt (starting at t and ending at $t + \delta t$) the participation ratio calculated at t remains applicable and constant over δt even though in reality its actual value may vary as the risky asset's volatility varies. This result is illustrated in Fig. 2.1 over two time intervals. The fact that the discretized participation ratio cannot compensate for all volatility movements means that hedging errors are introduced between the VTSs continuous dynamics previously derived and the discretized equivalents.

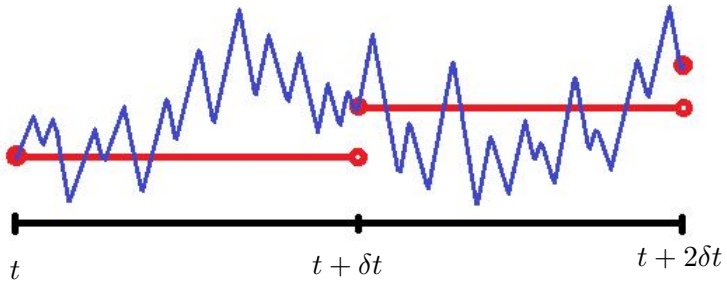

Fig. 2.1: A random walk of the participation ratio (ω_t) in continuous (blue) and discrete (red) time

In this dissertation one of the simplest time discretization schemes is employed, being the Euler scheme. Rouah (2013) provides several more complicated schemes that could be used to obtain more accurate discrete approximations. However since the aim of this dissertation is to show intuitively how SVMs can be used within a VTS framework, the Euler scheme's simplicity proves most illustrative.

For Euler, all integrals are evaluated using the left-point rule which assumes that values at time $t + \Delta t$ can be approximated with information known at t, as all time-dependant parameters become constant over Δt.

It is important to emphasise that the time interval (Δt) refers to the period over which returns are calculated i.e. how often the portfolio is rebalanced, whereas the time interval (δt) refers to the period over which the participation ratio remains constant i.e. how often model calibration occurs to obtain updated parameters.

Applying Euler's approximation scheme to equations (2.1) and (2.8) yields the

percentage returns of the risky asset

$$\begin{aligned} S_{t+\Delta t} &\approx S_t + \int_t^{t+\Delta t} \mu_u S_u \mathrm{d}u + \int_t^{t+\Delta t} \sqrt{V_u} S_u \mathrm{d}W_{1,u}^{\mathbb{P}} \\ &\approx S_t + \mu_t S_t \Delta t + \sqrt{V_t} S_t W_{1,\Delta t}^{\mathbb{P}} \\ &\approx S_t + \mu_t S_t \Delta t + \sqrt{V_t \Delta t} S_t Z_1 \\ R_S(t, t+\Delta t) = \frac{S_{t+\Delta t} - S_t}{S_t} &\approx \mu_t \Delta t + \sqrt{V_t \Delta t} Z_1 \end{aligned} \tag{2.12}$$

for the Heston model, and

$$\begin{aligned} S_{t+\Delta t} &\approx S_t + \int_t^{t+\Delta t} (\mu_u - \alpha) S_u \mathrm{d}u + \int_t^{t+\Delta t} \sqrt{V_u} S_u \mathrm{d}W_{1,u}^{\mathbb{P}} \\ &\quad + (J-1) \int_t^{t+\Delta t} S_u \mathrm{d}N_u \\ &\approx S_t + (\mu_t - \alpha) S_t \Delta t + \sqrt{V_t \Delta t} S_t Z_1 + (J-1) S_t N_{\Delta t} \\ R_S(t, t+\Delta t) = \frac{S_{t+\Delta t} - S_t}{S_t} &\approx (\mu_t - \alpha) \Delta t + \sqrt{V_t \Delta t} Z_1 + (e^{Z_3} - 1) N_{\Delta t} \end{aligned} \tag{2.13}$$

for the Bates model where $N_{\Delta t} \sim \mathcal{P}(\theta \Delta t)$, $\mathcal{P}$ being the Poisson distribution. Similarly applying Euler's scheme to the variance process equation (2.2) yields

$$\begin{aligned} V_{t+\Delta t} &\approx V_t + \int_t^{t+\Delta t} \kappa \left(\bar{V} - V_u\right) \mathrm{d}u + \int_t^{t+\Delta t} \sigma \sqrt{V_u} \mathrm{d}W_{2,u}^{\mathbb{P}} \\ &\approx V_t + \kappa \left(\bar{V} - V_t\right) \Delta t + \sigma \sqrt{V_t} W_{2,\Delta t}^{\mathbb{P}} \\ &\approx V_t + \kappa \left(\bar{V} - V_t\right) \Delta t + \sigma \sqrt{V_t \Delta t} \left(\rho Z_1 + \sqrt{1-\rho^2} Z_2\right) \end{aligned} \tag{2.14}$$

where Z_2 is standard normal random variate, independent of Z_1, both independent of Z_3 which was previously defined in section 2.1.2.

As stated in section 2.1.1, if the Feller condition holds equation (2.14) should never drop below zero. Rouah (2013) showed empirically that if enough paths of the Euler discretized variance process are simulated it is likely that at least one path will drop below zero even if the Feller condition holds. To overcome this problem he suggests implementing either a full truncation scheme or reflection scheme. Since the full truncation scheme will result in variance values equal to zero hence participation ratios equal to infinity, the reflection scheme is employed. Therefore V_t is replaced by $|V_t|$ everywhere in the discretization.

The discretized equivalent of the NACC risk-free asset can be modelled by assuming that the asset now accrues at a simple risk-free rate, over time interval Δt

$$\begin{aligned} B_{t+\Delta t} &= B_t (1 + r \Delta t) \\ \frac{B_{t+\Delta t} - B_t}{B_t} &= r \Delta t \end{aligned} \tag{2.15}$$

Substituting equations (2.12) (or (2.13)) and (2.15) into (1.2) yields the discrete price of the VTS at time $t + \Delta t$

$$X_{t+\Delta t} = X_t \left[1 + \omega_t R_S(t, t+\Delta t) + (1 - \omega_t) r \Delta t\right] \tag{2.16}$$

Now consider the value of the VTS at the following series of discrete time points $\{t; t+\Delta t; t+2\Delta t; ...; t+N\Delta t\}$

$$\begin{aligned}
\text{At time } t+2\Delta t \rightarrow X_{t+2\Delta t} &= X_{t+\Delta t}\left[1+\omega_{t+\Delta t}R_S\left(t+\Delta t, t+2\Delta t\right)+\left(1-\omega_{t+\Delta t}\right)r\Delta t\right] \\
&= X_t\left[1+\omega_t R_S\left(t, t+\Delta t\right)+\left(1-\omega_t\right)r\Delta t\right] \\
&\times\left[1+\omega_{t+\Delta t}R_S\left(t+\Delta t, t+2\Delta t\right)+\left(1-\omega_{t+\Delta t}\right)r\Delta t\right] \\
&\vdots \\
\text{At time } t+N\Delta t \rightarrow X_{t+N\Delta t} &= X_t\prod_{i=0}^{N-1}\left[1+\omega_m R_S\left(j,k\right)+\left(1-\omega_m\right)r\Delta t\right] \qquad (2.17) \\
\text{if } \delta t=\Delta t \rightarrow m &= \{t; t+\Delta t; t+2\Delta t; ...; t+(N-1)\Delta t\}=j \\
\text{if } \delta t=2\Delta t \rightarrow m &= \{t; t; t+2\Delta t; ...; t+(N-2)\Delta t\} \\
\text{if } \delta t=3\Delta t \rightarrow m &= \{t; t; t; ...; t+(N-3)\Delta t\} \\
&\vdots \\
\text{if } \delta t=N\Delta t \rightarrow m &= \{t; t; t; ...; t\}
\end{aligned}$$

where $j = t+i\Delta t$ and $k = t+(i+1)\Delta t$. The above result illustrates that when $\delta t = \Delta t$, a participation ratio is calculated at every interval that the portfolio is rebalanced therefore the calibration rate is equal to the rebalancing rate. For the case when $\delta t = 2\Delta t$, the participation ratio calculated at t remains applicable for two rebalancing intervals, only updating when m = 2. From the same intuition it can be shown that when $\delta t = N\Delta t$ the original participation ratio calculated at t remains applicable for all rebalancing periods.

It follows that the discrete dynamics obtained when $\delta t = \Delta t$ serve as the closest possible replication of the true continuous VTS. For the Heston model, the accuracy of this result can be analysed by calculating the hedging error between the portfolio price obtained from the closed-form solution defined in equation (2.6) and the discretized dynamics defined above. For the case of the Bates model there is no closed-form solution to test against the discretized dynamics, therefore the dynamics obtained when $\delta t = \Delta t$ serve as a continuous approximate for comparisons against less frequent calibration rates.

The hedging error can be defined as

$$HE = \sqrt{\frac{\Phi}{N-1}\sum_{i=0}^{N-1}\left[R_{\overline{X}}\left(j,k\right)-R_X\left(j,k\right)\right]^2} \qquad (2.18)$$

where $R_{\overline{X}}$ are the percentage returns of the VTS hedged using some calibration rate and R_X are the daily returns of the continuously hedged VTS (continuous proxy in the case of Bates). The $\sqrt{\Phi}$ term is used to annualize the hedging ratio, for instance in this dissertation, returns are calculated daily so Φ = 252 (assuming there are 252 business days in a year).

It must be noted that there is a degree of error introduced into equation (2.18) from the Euler discretization scheme when calculating the hedging error for X_t^H.

This error could be reduced by using a more accurate discretization scheme as previously mentioned, allowing the returns of equation (2.17) to tend closer to those of equation (2.6). This error is not relevant for X_t^B since there is no closed-form solution for the Bates driven VTS. Therefore the discretization that serves as the continuous approximation will itself contain a Euler discretization error.

2.3 Construction under the Risk-Neutral Measure

The VTS dynamics derived so far have been applicable under the real-world measure. Since the secondary objective of the VTS is to act as an underlying which options can be written on, it is necessary to establish the dynamics under the risk-neutral measure. To construct the VTS under the risk-neural measure it is assumed that there exists a market price of risk, previously defined as λ_s, for the risky asset. For the Heston model it is assumed that there exists only one possibility for λ_s implying the market is complete since there is only one source of noise present in the underlying dynamics. The same assumption of completeness cannot be made for the Bates model as the Poisson process acts as a second source of noise.

A VTS under the risk-neutral measure $\mathbb{Q}$ can be constructed from the real-world measure $\mathbb{P}$ via a Girsanov transformation whose kernel is negative the market price of risk. Mathematically this can be expressed as

$$\mathrm{d}W_{1,t}^{\mathbb{P}} = \mathrm{d}W_{1,t}^{\mathbb{Q}} - \lambda_s \mathrm{d}t \tag{2.19}$$

for the Heston modelled VTS and

$$\mathrm{d}W_{1,t}^{\mathbb{P}} = \mathrm{d}W_{1,t}^{\mathbb{Q}} - \left(\frac{\sqrt{V_t + \beta}}{\sqrt{V_t}}\right) \lambda_s \mathrm{d}t \tag{2.20}$$

for the Bates modelled VTS, where $\mathrm{d}W_{1,t}^{\mathbb{Q}}$ is a Brownian motion process applicable for the risk-neutral measure $\mathbb{Q}$. Equation (2.20) holds so long as the variance process never reaches zero. Substituting equations (2.19) and (2.20) into (2.1) and (2.8) and solving yields the risk-neutral dynamics of the risky asset which can be defined as

$$\frac{\mathrm{d}S_t}{S_t} = r\mathrm{d}t + \sqrt{V_t}\mathrm{d}W_{1,t}^{\mathbb{Q}} \tag{2.21}$$

for the Heston model, and

$$\frac{\mathrm{d}S_t}{S_t} = (r - \alpha)\,\mathrm{d}t + \sqrt{V_t}\mathrm{d}W_{1,t}^{\mathbb{Q}} + (J - 1)\mathrm{d}N_t \tag{2.22}$$

for the Bates model.

To obtain the risk-neutral dynamics of the CIR variance process, a similar Girsanov transformation needs to be implemented. The difference comes in that previously the market price of risk was used to adjust the drift, now a measure known as the volatility risk premium is used. Mathematically this was defined by Rouah (2013) as

$$\mathrm{d}W_{2,t}^{\mathbb{P}} = \mathrm{d}W_{2,t}^{\mathbb{Q}} - \left(\frac{\sqrt{V_t}}{\sigma}\right) \lambda_v \mathrm{d}t \tag{2.23}$$

where λ_v is the volatility risk premium embedded in the quoted volatility surface. Equation (2.23) holds so long as the volatility of variance is strictly greater than zero. Such a constraint could be easily met by setting up the appropriate calibration boundaries. Volatility risk premium refers to the observation that implied volatilities quoted in the market usually exceed the realized volatility of the underlying asset as defined by Ge (2016). The main causal mechanism for such occurrences is the risk averse nature of most investors. As previously mentioned when stock prices are low, volatility is usually quoted high to favour option writers. Consequently one could expect volatility to be slightly inflated in the risk-neutral world when compared to what the underlying asset is actually experiencing in reality. This inflation is what λ_v represents in equation (2.23).

Substituting equation (2.23) into (2.2) gives the risk-neutral dynamics of the variance process

$$\mathrm{d}V_t = \tilde{\kappa}\left(\Upsilon - V_t\right)\mathrm{d}t + \sigma\sqrt{V_t}\mathrm{d}W^{\mathbb{Q}}_{2,t} \tag{2.24}$$

where $\tilde{\kappa} = \kappa + \lambda_v$ and $\Upsilon = \frac{\kappa\bar{V}}{\kappa+\lambda_v}$. Note when $\lambda_v = 0 \Rightarrow \tilde{\kappa} = \kappa$ and $\Upsilon = \bar{V}$. When calibrating model parameters from quoted surfaces the parameters obtained already contain λ_v embedded in them, therefore it is not necessary to estimate the volatility risk premium unless one wishes to physically harvest it. Since this is not the goal of this dissertation it is assumed that $\lambda_v = 0$ throughout, therefore the risk-neutral model parameters are analogous to the real-world. For more information on volatility risk premium harvesting techniques consult Ge (2016) or Israelov and Klein (2016).

With the risk-neutral dynamics defined in equations (2.21), (2.22) and (2.24), a similar procedure conducted for the real-world measure can be performed to obtain the distribution properties of X^H_t and X^B_t, as well as their discretized dynamics and related hedging errors.

It can be easily shown that the return distribution of X^H_t, under the risk-neutral measure, is equal to that of an asset following GBM where the drift is equal to the risk-free rate. It follows that the BS formulae can be used to calculate the value of options written on the VTS where the BS constant volatility is equal to σ_{target}. This would be the ideal case of option pricing using the VTS where the participation ratio is continuously calculated leading to instantaneous portfolio hedging. Of course in reality, continuous hedging is not possible and there will be a degree of error introduced by the calibration rate.

For X^B_t the prevalence of jumps will determine how closely the VTS dynamics resemble GBM and therefore how close an option priced on X^B_t tends towards the BS price.

Chapter 3

Implementation of the VTS

The purpose of this chapter is to illustrate how the VTS could be constructed practically using stochastic volatility estimates, as well as how implied volatilities (obtained from vanilla options written on the VTS) behave.

To obtain these estimates model parameters need to be calibrated from implied volatility surfaces observed in the market. Therefore it is imperative that there exists a time series of option data available for implementation. The underlying of these options serves as the risky asset for the VTS.

As previously mentioned volatility exhibits several stylised facts which make it a prime candidate for modelling using SVMs, however there is an additional advantage which needs to be emphasised. When using statistical models such as an Exponentially Weighted Moving Average (EWMA) or GARCH models to estimate volatility there is a degree of dependence on historical observations. The assumption made here is that to some degree, the current level of volatility measured in the market is influenced by past information. There are many reported cases where this assumption fails, typically around periods of financial crisis where sharp downward directed jumps are observed which could not have been foreseen by historical volatility preceding the crisis, as explained by Hull *et al.* (2009). The benefit of using quoted option data to obtain sequential volatility estimates is that each volatility surface can be interpreted as the market's overall perception of volatility for that specific point in time, looking forward, irrespective of past information. Such a perception proves useful considering the fact that the driving force behind the price fluctuations of the underlying asset can be attributed to a similar market, hence the opinion on the trading value of an option offers some insight as to where the underlying price is expected to be.

Fig. 3.1 compares the VIX index to the annualized, 21 day historical volatility of the S&P500 index. The VIX index is calculated[1] by taking a weighted average of implied volatilities for several call and put options written on the S&P500 index, therefore it represents a risk measure obtained from the information available from implied volatility surfaces. It can be seen that a strong correlation exists between the the two curves, especially during periods of high historical volatility. The degree that these curves deviate by can be attributed to the volatility risk premium (λ_v) that exists in the market.

[1] For more information on how the VIX index is calculated, visit http://www.cboe.com/micro/vix/vixwhite.pdf

Fig. 3.1: Bloomberg snapshot showing the relation between the VIX index and S&P500's historical volatility from 31 December 2007 to 31 December 2017

A trader looking to implement the VTS would have to decide on the rate at which the portfolio is rebalanced, as well as the rate at which calibration takes place. These rates would depend fully on the availability of data, and the willingness of the trader to act upon that data considering there may be costs incurred for doing such. Again consider the following series of discrete time points $\{t, t+\Delta t, t+2\Delta t, ..., t+N\Delta t\}$ which represent the points in time where market data becomes available, with t being today and Δt being 1 day. The underlying risky asset's price can be denoted as $\{S_t, S_{t+\Delta t}, S_{t+2\Delta t}, ..., S_{t+N\Delta t}\}$ with its quoted volatility surfaces being denoted as $\{\sigma_t(K,T), \sigma_{t+\Delta t}(K,T), \sigma_{t+2\Delta t}(K,T), ..., \sigma_{t+N\Delta t}(K,T)\}$. Here K represents the strike price, whereas T represents the maturity of the option. The risk-free asset's price can be denoted as $\{B_t, B_{t+\Delta t}, B_{t+2\Delta t}, ..., B_{t+N\Delta t}\}$ which is simply a chosen notional accruing at the overnight risk-free rate.

At time t the trader uses $\sigma_t(K,T)$ to estimate the values of V_0, μ_J, σ_J and θ (last 3 parameters used to calculate β in the case of Bates). The trader now has an estimate of the level of volatility expected by the market. With this volatility estimate, ω_t can be calculated to determine the weighting that must be allocated to the risky and risk-free asset. If the trader wishes to rebalance and recalibrate daily then at time $t+\Delta t$ the surface $\sigma_{t+\Delta t}(K,T)$ would be used to obtain new values of V_0, μ_J, σ_J and θ and so the process would continue daily over time.

3.1 Simulation & Hedging Error Testing

To successfully back test the VTS one would need a sufficiently long time series of historical volatility surfaces to obtain sequential volatility estimates. The result is a tedious exercise of surface calibration which falls out of the scope of this dissertation. Therefore instead of using market data, all the necessary data will be simulated using the framework defined in chapter 2. That is the risky asset can be simulated using equations (2.12) and (2.13), the risk-free asset can be simulated using equation (2.15) and the volatility of the risky asset can be simulated using equation (2.14). To simulate the time-evolving model parameters that would be obtained from periodic calibration, each parameter is allowed to accrue by some degree of weighted noise. This can be expressed as

$$x_{t+\delta t} = x_t \pm 10\text{bps} \times \text{randn}()$$

where $x_t = \{\kappa, \bar{V}, \sigma, \rho, \mu_J, \sigma_J, \theta\}$ and the noise term is denoted by a standard normal random variable whose magnitude is adjusted by 10 basis points. Boundary conditions are set for each parameter to ensure that noise term never drives a parameter to an infeasible level. The boundary conditions are defined as

$$\begin{aligned} \{\kappa, \sigma\} &= (0, 5] \\ \{\bar{V}, \sigma_J, \theta\} &= (0, 1] \\ \{\rho, \mu_J\} &= [-1, 1] \end{aligned}$$

If at any point in time $x_{t+\delta t}$ reaches a boundary the noise term is removed from x_t, else it is added. These boundary conditions are synonymous with the parameter limits that would have been set during calibration. The initial model parameters used for the base simulation are given in Tab. 3.1

V_0	κ	$\bar{V}$	σ	ρ	μ_J	σ_J	θ
0.01	1.7	0.0225	0.3	-0.2	-0.4	0.1	0.001

$X_0 = S_0 = B_0$	λ_s	r	t_0	$\delta t = \Delta t$	t_{end}
100	0.1	0.032	0	1/252	10

Tab. 3.1: Initial Heston and Bates model parameters used for base simulation

The simulated paths of the risky asset (with its volatility), risk-free asset and VTS (with $\delta t = \Delta$t) are given in Fig. 3.2 for σ_{target} = 5%, 12% and 30%. It must be noted that the curves of the risky and risky-free asset remain constant therefore the scale of the price axis depends purely on the performance of the VTS for a particular σ_{target}. The seed was fixed to ensure all simulations experience the same random walk. It is evident for both the Heston and Bates simulations that when the targeted volatility $\ll$ measured volatility, the VTS evolution partially resembles the risk-free asset with little noise introduced into its dynamics. When the targeted volatility $\approx$ measured volatility, the VTS resembles the risky asset except that it is less affected by periods of high volatility and outperforms the risky asset for both models.

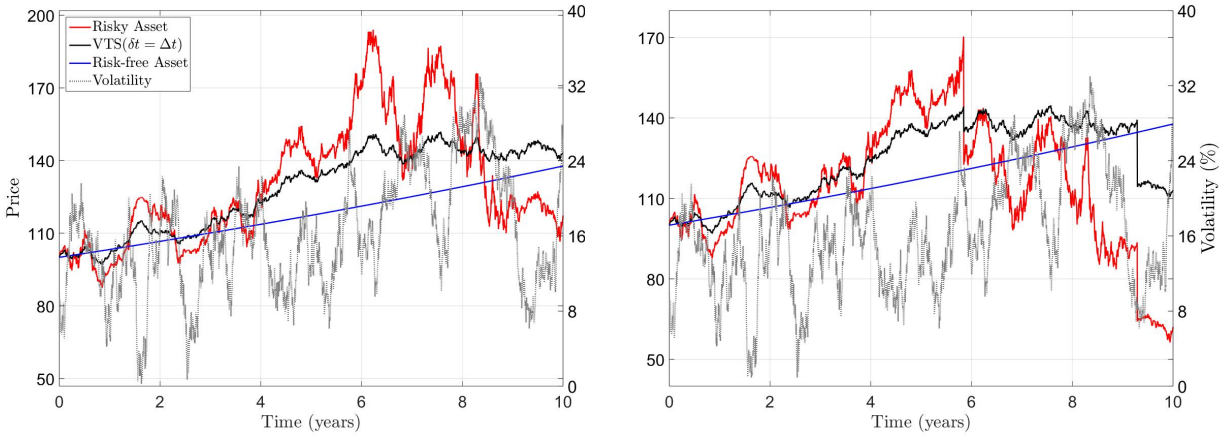

(a) Heston simulation with $\sigma_{target} = 5\%$ **(b)** Bates simulation with $\sigma_{target} = 5\%$

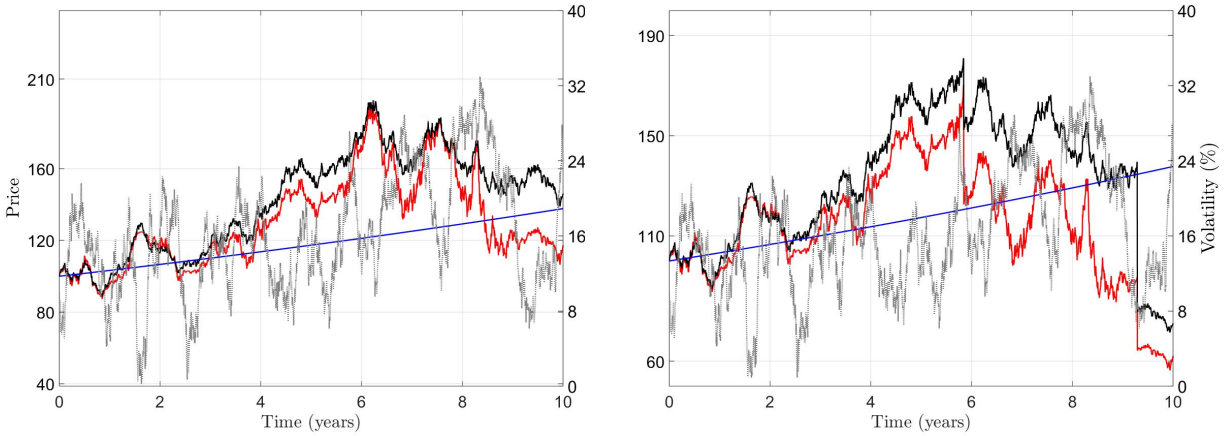

(c) Heston simulation with $\sigma_{target} = 12\%$ **(d)** Bates simulation with $\sigma_{target} = 12\%$

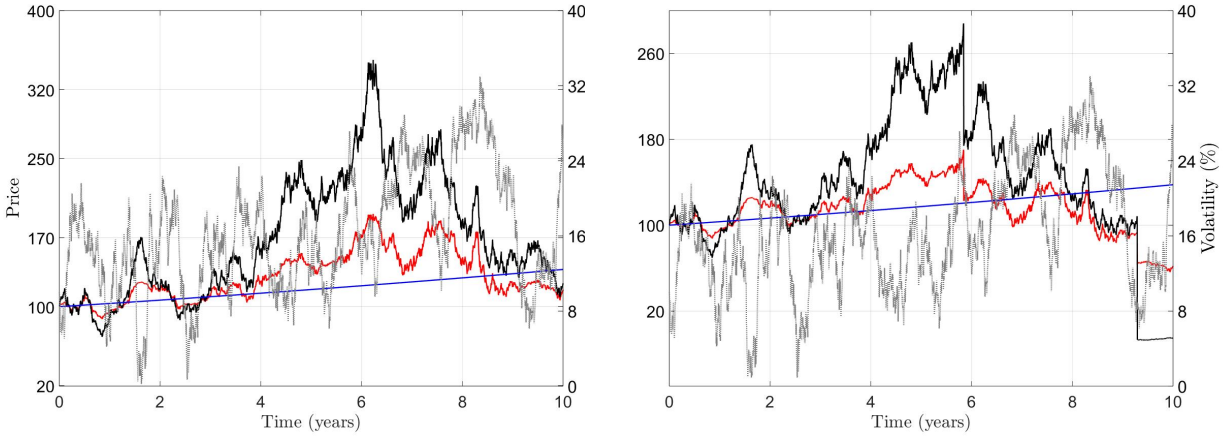

(e) Heston simulation with $\sigma_{target} = 30\%$ **(f)** Bates simulation with $\sigma_{target} = 30\%$

Fig. 3.2: Simulation of Heston and Bates modelled risky asset, risk-free asset, volatilities and VTS for σ_{target} =5%,12% and 30% and $\Delta t = \delta t$

Finally, for the case when the targeted volatility $\gg$ measured volatility, the VTS far outperforms the risky asset when the measured levels of volatility are low. From Fig. 3.2 it might seem obvious that one would want to choose a high σ_{target} to optimize returns but there is an important consequence to owning more of the risky asset which is observable by the way in which the portfolio reacts to jumps.

The first jump occurs roughly at year 6. At this point volatility peaks at a high of 28% resulting in a low participation ratio and a small allocation of the risky asset to the VTS, therefore the jump does not heavily influence the portfolio. The opposite occurs at the second jump which is located just after year 9. At this point volatility reaches a relative low of 8% resulting in a high participation and large allocation of the risky asset to the VTS. Consequently the jump heavily effects the VTS.

For the above base simulation it was assumed that the portfolio rebalancing rate was equal to the calibration rate i.e. $\Delta t = \delta t$. Now the VTS is simulated for varying calibration rates to test how a practically constructed VTS would differ from the continuous mathematical constructions defined in equations (2.5) and (2.9). It was shown that for X_t^H a closed-form solution exists, defined in equation (2.6), which serves as the continuous benchmark. X_t^B has no closed-form solution therefore the solution when a calibration rate of $\Delta t = \delta t$ serves as an approximate continuous benchmark. The test is performed by calculating the hedging error between the continuous returns and the returns obtained from some calibration rate, as defined by equation (2.18). The test results are summarized in Tab. 3.2. For all simulations the portfolio was rebalanced at a daily rate for 10 years.

	Daily $\delta t = \Delta t$	**Weekly** $\delta t = 5\Delta t$	**Monthly** $\delta t = 20\Delta t$	**Quarterly** $\delta t = 60\Delta t$
σ_{target} = 5%				
Heston Hedging Error (%)	0.0104	0.9527	2.6590	8.3498
Bates Hedging Error (%)	0	0.6854	2.3968	6.3859
σ_{target} = 12%				
Heston Hedging Error (%)	0.0600	2.2872	6.3824	20.0388
Bates Hedging Error (%)	0	1.6450	5.7523	15.3260
σ_{target} = 30%				
Heston Hedging Error (%)	0.3751	5.7282	15.964	50.093
Bates Hedging Error (%)	0	4.1126	14.381	38.315

Tab. 3.2: Heston and Bates hedging errors for varying calibration rates and target volatilities

The annualized hedging error was shown to increase as the target volatility increases which is intuitive as the more volatility is allowed to drive the VTS the less likely the VTS will obtain the constant volatility characteristics synonymous with

the continuous benchmark, since the source of the volatility is non-constant itself. It is also shown, for both models, that the hedging error increases drastically as the calibration rate decreases from daily to quarterly. This result is expected and ties back to the proposition that continuous hedging can be approximated by a higher hedging frequency.

It is also observed that the hedging error of X_t^H is greater than that of X_t^B. This is again expected as an additional Euler discretization error is introduced into the total hedging error for X_t^H. As stated previously such a hedging error is not present for X_t^B as there is no continuous benchmark to compare against. One should interpret these results with this in mind. In reality the hedging error of X_t^B would be greater than that of X_t^H, depending on the magnitude of jumps incurred.

3.2 Extracting Implied Volatilities from the VTS

In chapter 2.3 the risk-neutral dynamics of the risky asset were defined. Using these dynamics, the risk-neutral equivalents of equation (2.17) can be used to simulate path-wise estimates of a VTS whose drift is equal to the risk-free rate. Consequently options written on the VTS can be priced using standard Monte Carlo pricing methodology.

The ability to price under the risk-neutral measure is a direct consequence of using implied volatility estimates (over historical) to calculate the participation ratio. If historical estimates were used one would be left with an unknown, real-world distribution for the returns of the VTS. Implied estimates on the other hand can be used to obtain SVM parameters through calibration, where the dynamics have a known distribution. This makes option pricing possible.

Stutzer (1996) provides a non-parametric approach to option pricing where a risk-neutral distribution is approximated from an empirical distribution obtained from an underlying's historical price. Incorporating such an approach is advantageous because it does not need observable market implied volatilities as all option pricing information is obtained directly from the underlying's historical price. The major disadvantage of this approach is that the risk-neutral distribution calculated is of course an approximation and can differ greatly from the expected risk-neutral distribution. For this reason it is assumed that at all future calibration times a surface of liquid options exist that a SVM can be fitted to.

Since options are quoted in volatility, it proves useful to display these option prices in volatility magnitudes. The implied volatility that would yield an option price can be approximated using the standard BS formulae. The following steps explain how these implied volatility surfaces can be obtained from options written on the VTS:

1. At time t model parameters are calibrated from implied volatility surface $\sigma_t(K, T)$.

2. The risk-neutral equivalents of equations (2.12), (2.13) and (2.14) can be used to simulate n paths of the risky asset with its volatility. This is accomplished

by generating n independent and identically distributed (i.i.d) normal random variates for Z_1, Z_2 and Z_3 and Poisson random variates for $N_{\Delta t}$ at each time increment Δt.

3. The aforementioned equations can be used (with the deterministic risk-free asset simulated using equation (2.15)) to determine n many risk-neutral equivalents of X_T^n where $T = t + N\Delta t$ in equation (2.17).

4. The price of a European option written on the VTS can then be calculated by taking the discounted average of all n pay-offs at maturity. This can be expressed mathematically as

$$\begin{aligned} O(K,T) &= \exp(-rT)\mathbb{E}^{\mathbb{Q}}\left[f\left(X_T^n\right)\right] \\ &= \frac{\exp(-rT)}{n}\sum_{p=1}^{n} f\left(X_T^p\right) \end{aligned} \tag{3.1}$$

where $f\left(X_T^n\right) = \left(X_T^n - K\right)^+$ is the pay-off for a call option and $\left(K - X_T^n\right)^+$ is the pay-off for a put option.

5. The BS formula can then be used to approximate the implied volatility that would return the option price calculated from equation (3.1), since the only unknown in the proceeding BS formula is $\tilde{\sigma}_t(K,T)$

$$O(K,T) = \eta X_0 \mathcal{N}(\eta d_1) - \eta \exp(-rT) K \mathcal{N}(\eta d_2) \tag{3.2}$$

where $d_1 = \frac{\ln\left(\frac{X_0}{K}\right) + \left(r + \frac{1}{2}\tilde{\sigma}_t(K,T)^2\right)T}{\tilde{\sigma}_t(K,T)\sqrt{T}}$, $d_2 = d_1 - \tilde{\sigma}_t(K,T)\sqrt{T}$ and η = 1 for a call option and -1 for a put option.

$\mathcal{N}()$ is the cumulative distribution function of the standard normal distribution.

6. The value of $\tilde{\sigma}_t(K,T)$ that minimizes the difference between equations (3.1) and (3.2) is the approximated implied volatility of an option written on the VTS. Matlab's fzero function was used to determine this approximate. Fzero is a non-linear root finding function which tries to find a point x where $fun(x) = 0$. The solution is where $fun(x)$ changes sign.

It must be noted that if S_T was used instead of X_T in equation (3.1) and S_0 instead of X_0 in equation (3.2), the volatility surface $\sigma_t(K,T)$ would be returned for the Heston and Bates modelled risky assets.

While performing step 6 of the above analysis it was observed that most of the extreme in-the-money (ITM) and OTM implied volatilities failed to converge for short maturity dates. That is, fzero failed to find an approximate implied volatility and returned the initial guess input. This result was more prevalent for ITM cases. The intuition behind such a result is that since it is very unlikely that these options will end up OTM in the short time till expiry they would demand a very high price, implying a quoted implied volatility that would tend towards infinity. To overcome

this problem only OTM options were priced, hence if $X_0 \leq K$ call options were considered and if $X_0 > K$ put options were considered.

After implementing such a constraint it was observed that a few of the extreme OTM implied volatilities failed to converge. The intuition behind this result is that since it is very unlikely that these options will end up ITM in the short time till maturity, their prices would tend towards zero implying an implied volatility close to zero. Since the BS formula does not allow for an implied volatility of zero ($\tilde{\sigma}_t(K,T)$ appears in the denominator of d_1) fzero ends up returning the initial guess input. To overcome this issue a Vega-weighted second order polynomial was fitted to the available set of implied volatilities using least squares regression. Mathematically the Vega-weighted polynomial can be expressed as

$$y = a_2 x^2 + a_1 x + a_0$$

where $y = \tilde{\sigma}_t(K,T)^2 T$ and $x = \log\left(\frac{K}{X_0}\right)$. Chapra and Canale (1998) prove that the constants a_0, a_1 and a_2 can be represented by the set of linear equations

$$\begin{aligned}
(L)\, a_0 + \left(\sum_{i=1}^{L} x_i\right) a_1 + \left(\sum_{i=1}^{L} x_i^2\right) a_2 &= \sum_{i=1}^{L} y_i \\
\left(\sum_{i=1}^{L} x_i\right) a_0 + \left(\sum_{i=1}^{L} x_i^2\right) a_1 + \left(\sum_{i=1}^{L} x_i^3\right) a_2 &= \sum_{i=1}^{L} x_i y_i \\
\left(\sum_{i=1}^{L} x_i^2\right) a_0 + \left(\sum_{i=1}^{L} x_i^3\right) a_1 + \left(\sum_{i=1}^{L} x_i^4\right) a_2 &= \sum_{i=1}^{L} x_i^2 y_i
\end{aligned}$$

where i refers to the vector of strikes that have successfully calculated implied volatilities without need of the Vega-weighted correction and L is the length of that vector. To anchor the polynomial's minimum value to the ATM implied volatility, a_0 is set equal to $\tilde{\sigma}_t(X_0,T)^2 T$. The extreme OTM implied volatilities can then be approximated by setting $\tilde{\sigma}_t(K,T) = \sqrt{\frac{y}{T}}$ with a_0, a_1 and a_2 known for a specific T and K equal to the strikes that were unsuccessful in calculating implied volatilities. This method of using polynomials to fit implied volatility surfaces was first explored by Dumas *et al.* (1998).

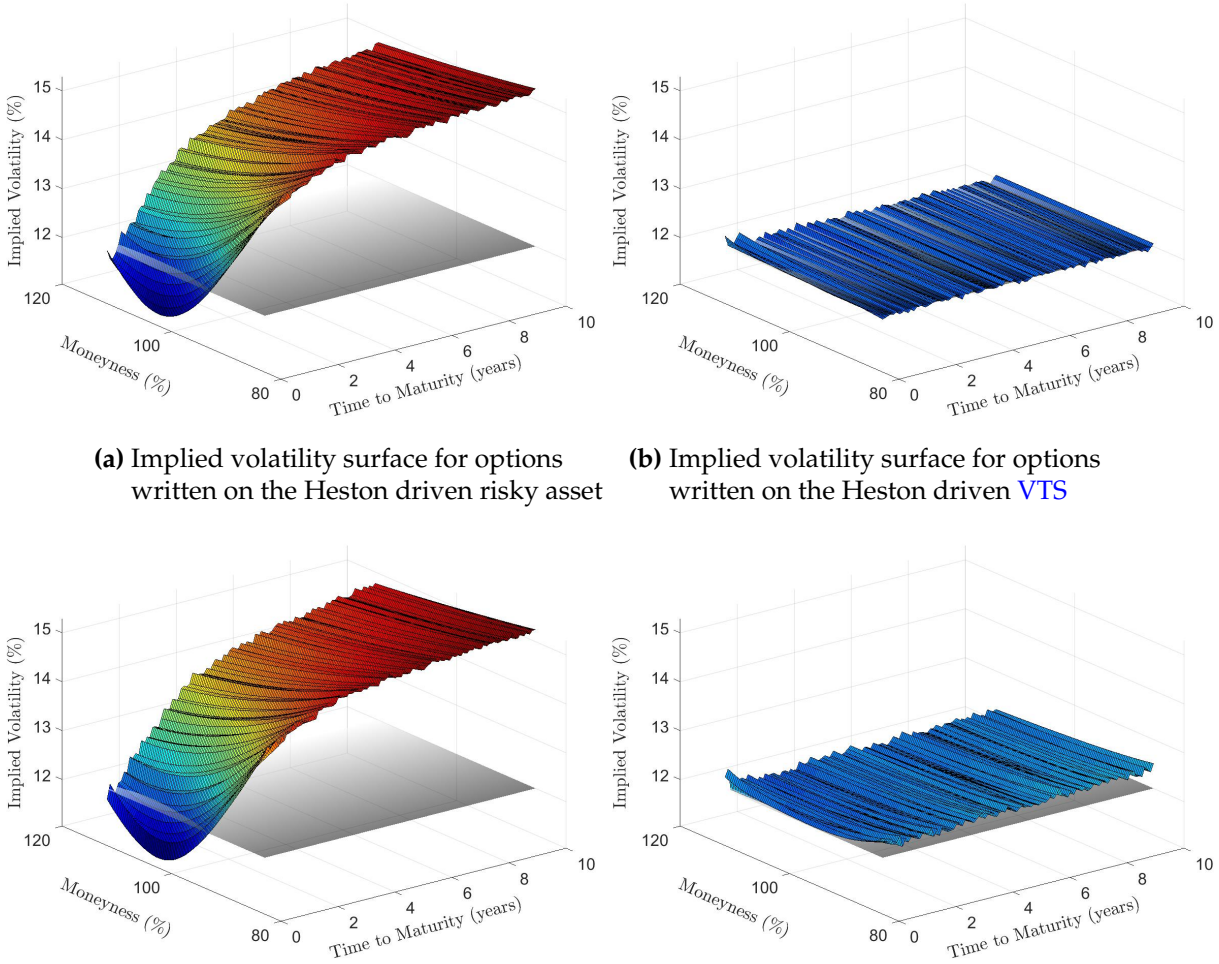

(a) Implied volatility surface for options written on the Heston driven risky asset

(b) Implied volatility surface for options written on the Heston driven VTS

(c) Implied volatility surface for options written on the Bates driven risky asset

(d) Implied volatility surface for options written on the Bates driven VTS

Fig. 3.3: Implied volatility surfaces obtained from a MCS with n = 100 000, σ_{target} = 12% and $\delta t = \Delta t$. The implied volatilities are plotted against moneyness levels of 85% to 115% and maturities of 1.2 months out to 10 years at intervals of 1.2 months

Fig. 3.3 illustrates the implied volatility surfaces obtained from the aforementioned methodology. The model parameters defined in Tab. 3.1 were used for the MCS. It must be noted that since $X_0 = S_0$, the measure of moneyness for the risky asset is synonymous with that of the VTS and is defined as the spot price divided by the strike price. It must also be noted that implied volatilities where moneyness $\in [85\%, 100\%)$ were obtained from put options whereas implied volatilities where moneyness $\in [100\%, 115\%]$ were obtained from call options. It was observed that for options written on the VTS, both X_t^H and X_t^B return relatively flat implied volatility surfaces.

An exception is observed for implied volatilities obtained from OTM puts writ-

ten on X_t^B. This result can be directly attributed to the Bates model parameters that drive the Poisson process. An option that is very far OTM with a short time to maturity can only attain significant value if rapid movement of its underlying is likely, that is for a call rapid positive movement and for a put rapid negative movement. By setting $\mu_J < 0$, sharp downward directed jumps are included in the modelling framework therefore it becomes possible for a put option that is far OTM to end up ITM. The result is obvious for all far OTM puts where their volatilities converge to a level just above $\sigma_{target} = 12\%$, which can be understood as the inflated price of these OTM puts whose price is conditional on the possibility of a sharp downward directed jump.

Other than this OTM put observation, the VTS successfully manages to return a flat implied volatility surface from a risky asset which posses a large degree of smile and term-structure embedded in its underlying price.

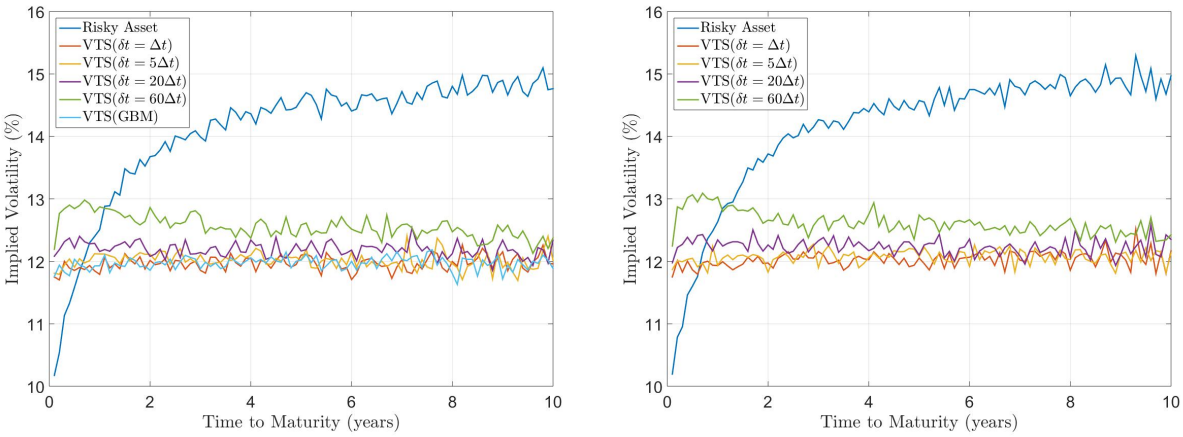

(a) Implied volatilities ATM options written on the Heston driven risky asset and VTS

(b) Implied volatilities ATM options written on the Bates driven risky asset and VTS

Fig. 3.4: At-the-money implied volatility for the risky asset and VTS obtained from a MCS with n = 100 000, σ_{target} = 12% and varying calibration rates, plotted against maturities of 1.2 months out to 10 years at intervals of 1.2 months

The implied volatility surfaces illustrated in Fig. 3.3 assumes that the underlying VTSs volatility is adjusted on a daily basis. In reality however the underlying VTS might be adjusted by some other frequency. Fig. 3.4 considers the implied volatility of ATM calls priced using the risky asset as well as VTSs (generated by using different calibration rates) as the underlying. The rates used are consistent with those given in Tab. 3.2, that is daily, weekly, monthly and quarterly calibration rates are used. For X_t^H the GBM dynamics provided by the risk-neutral equivalent of equation 2.6 are used to obtain a closed-form ATM implied volatility.

It is observed that there is little difference between the behaviour of the volatilities obtained from the two models. Furthermore the closed-form solution produces very similar results to the implied volatilities obtained from daily and weekly cal-

ibration. Observable differences begin to occur at monthly calibration, where it is now obvious that the volatility is inflated above σ_{target} for all maturity dates. For a quarterly rate the volatility seems heavily inflated for short maturity dates but begins to reduce as maturity increases. The above result concludes that for both models, there is little difference between daily and weekly calibration as both rates converge relatively well towards the closed-form solution and $\sigma_{target} = 12\%$.

Chapter 4

Pricing Path-Dependent Options Written on the VTS

Retirement annuity products offered by life insurers are characterized by two main features. Firstly their investment horizons are far longer than those offered by banking institutions. Here liquidity is usually limited to 2 years whereas the minimum tenor period of a retirement product is 10 years. Secondly the pay-off of the product should have a guarantee attached to it, ensuring that the buyer receives the agreed settlement regardless of the state of the world at maturity.

A lookback option is a form of path-dependent product that can achieve this level of guarantee. They do so by settling on some minimum (maximum) underlying value, registered over the life of the contract. At maturity the holder can essentially "lookback" and select the most convenient price that occurred during this period. For example purchasing a floating strike lookback put offers the buyer the opportunity to own the asset at its highest price. The cost of such certainty comes in the form of monthly premium payments, determined from the fair value of the option at inception.

Due to the non-Gaussian nature of returns, coupled with non-constant volatility features over long investment horizons, pricing and hedging of path-dependent options become problematic. To compensate for unknown risk factors, premiums are often inflated resulting in an expensive product for the purchaser. There are methods of obtaining closed-form prices, but all require returns to be normally distributed and cannot account for fluctuating volatility levels. Some of these methods include BS, finite difference and binomial trees. Alternative models that treat return distributions as non-normal do a better job of capturing unknown risk factors and are the industry standard at the expense of increased model complexity.

The VTS offers an underlying on which lookback options can be written on, due to its adjusted returns and constant volatility properties. Consequently the problem of pricing becomes a simple exercise of implementing a BS closed-form solution. To test the accuracy of this statement the BS price of a floating strike lookback put will be compared to a numerically approximated price, obtained from an error corrected MCS. For the BS price the input volatility is equal to σ_{target}. For the MCS, daily and weekly calibration will be used to interpret how the frequency at which volatility is updated impacts the accuracy of the numerically obtained price.

4.1 Floating Strike Lookback Put Option

The continuously sampled pay-off of a lookback put with floating strike is given by

$$f(X^*, X_T) = (X^* - X_T)^+ = X^* - X_T$$

where

$$X^* = \sup_{0 \leq u \leq T} X_u$$

implying $X_T \leq X^*$. The value of the option is equal to the discounted, expected value of this pay-off under the risk-neutral measure

$$LBP = \exp(-rT) \cdot \mathbb{E}^{\mathbb{Q}}[f(X^*, X_T)] = \exp(-rT) \cdot \mathbb{E}^{\mathbb{Q}}[X^* - X_T] \tag{4.1}$$

Conze (1991) gives a BS closed-form solution for a lookback put with floating strike written on an underlying that follows GBM.

$$\begin{aligned} LBP = X_0 \Big[& \exp(-rT)\mathcal{N}(-d + \sigma_{target}\sqrt{T}) - \mathcal{N}(-d) \\ & + \exp(-rT)\frac{\sigma^2_{target}}{2r}\left(-(1)^{\frac{-2r}{\sigma^2_{target}}} \cdot \mathcal{N}\left(d - \frac{2r\sqrt{T}}{\sigma_{target}}\right) + \exp(rT)\mathcal{N}(d)\right)\Big] \end{aligned} \tag{4.2}$$

where

$$d = \frac{\left(r + \frac{1}{2}\sigma^2_{target}\right)T}{\sigma_{target}\sqrt{T}}$$

and $\mathcal{N}()$ is the standard cumulative normal distribution function.

It is evident that equation (4.1) can be used to obtain a numerical price for the lookback put by simulating n many paths of X_T and X^*. The question thus arises, what is the best method of estimating X^*? Since it is impossible to obtain a continuous estimate, a trivial result would be to let X^* equal

$$X^* = \max(X_t, X_{t+\Delta t}, ..., X_{t+N\Delta t})$$

introducing a discretization error into the calculation. Anderson and Brotherton-Ratcliffe (1996) proposed a method that reduces the discretization error and improves the Monte Carlo estimate for a lookback options price. They achieved this by calculating the probability distribution of X^* over each time interval Δt, conditional on the change of $X_{\Delta t}$. Their calculation was based on the fact that these probability distributions are directly related to the probability distribution of the time of the first hit of a barrier, where the barrier in this case is a variable representing the maximum value. The conditional probability distribution defined above can be expressed mathematically as

$$\mathbb{P}(X^*_t \geq \ln X^*_t | X_t \in \ln X_{\Delta t}) = \exp\left(\frac{-2(\ln X^*_t - \ln X_{t+\Delta t})(\ln X^*_t - \ln X_t)}{\sigma^2_{target}\Delta t}\right) \tag{4.3}$$

The proof of equation (4.3) is given in Appendix A. In its derivation it is assumed that X_t follows GBM. This is not an unfair assumption to make on account of the

fact that if the VTS is successfully implemented, the dynamics should closely resemble that of GBM. To approximate a value for X_t^*, equation (4.3) is set equal to a standard uniformly distributed random variate u. After some algebra the result is equal to

$$\ln^2 X_t^* - \ln\left(X_t \cdot X_{t+\Delta t}\right) \ln X_t^* + \ln X_t \cdot \ln X_{t+\Delta t} + \frac{1}{2}\sigma_{target}^2 \Delta t \ln u = 0$$

Solving the above quadratic yields

$$X_t^* = \exp\left[\frac{1}{2}\left(\ln X_t + \ln X_{t+\Delta t} + \sqrt{\left(\ln X_t - \ln X_{t+\Delta t}\right)^2 - 2\sigma_{target}^2 \Delta t \ln u}\right)\right] \tag{4.4}$$

Therefore the true maximum over all time intervals can then be approximated as

$$\overline{X^*} = \max\left(X_t^*, X_{t+\Delta t}^*, ..., X_{t+N\Delta t}^*\right)$$

Equation (4.1) becomes

$$LBP = \exp\left(-rT\right) \cdot \mathbb{E}^{\mathbb{Q}}\left[\overline{X^*} - X_T\right] \tag{4.5}$$

It must be noted that if the alternative solution to quadratic was taken, then a negative sign would appear in front of the square root in equation (4.4) and it now returns the minimum value over the interval, which could be used to calculate a floating strike lookback call or a fixed strike lookback put.

4.2 Pricing Results

Equation (4.5) was used to perform an error corrected MCS whereby a floating strike lookback put was priced. The Heston and Bates SVM parameters defined in Tab. 3.1 were used to simulate the paths, where daily and weekly calibration rates were used. The target volatilities illustrated in Fig. 3.2 are again used to price the lookback put, that is $\sigma_{target} = 5\%, 12\%$ and 30%. The sample size used for the simulation was increased from 1 000 to 50 000 at intervals of 1 000 so that standard deviation bounds could be plotted around the price obtained from equation (4.5). The results are illustrated in Fig. 4.1. The solid red line represents the BS price obtained from using equation (4.2). The black dots are the daily calibrated prices, whereas the grey dots are the weekly calibrated prices. The dashed lines that encompass the respective dots are the standard deviation bounds.

It was shown that for all three target volatilities, the daily calibrated Heston driven VTS returns a price that converges onto the BS closed-form price. For the daily calibrated Bates driven VTS the price converges to a value slightly above the closed-form solution. Both models show that weekly calibration fails to converge to a price near the closed-form price. It also becomes apparent that this price fails to converge as sample size is increased, which is reflected by the way in which some high sample sizes experience sharp peaks in their standard deviation bounds. This result is especially prevalent for weekly calibration when $\sigma_{target} = 30\%$.

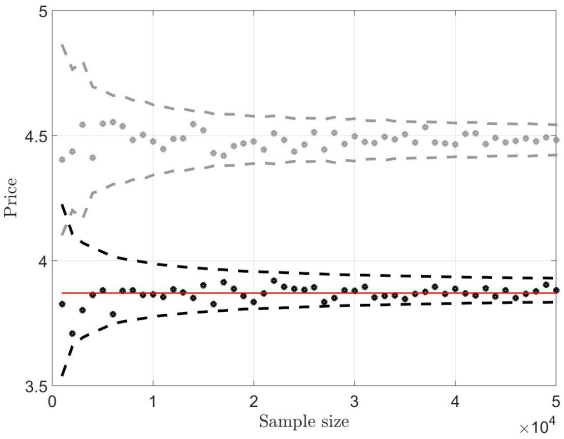

(a) Heston driven VTS used as an underlying with $\sigma_{target} = 5\%$

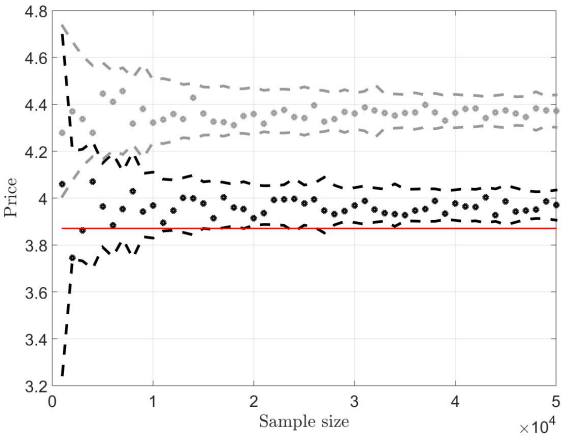

(b) Bates driven VTS used as an underlying with $\sigma_{target} = 5\%$

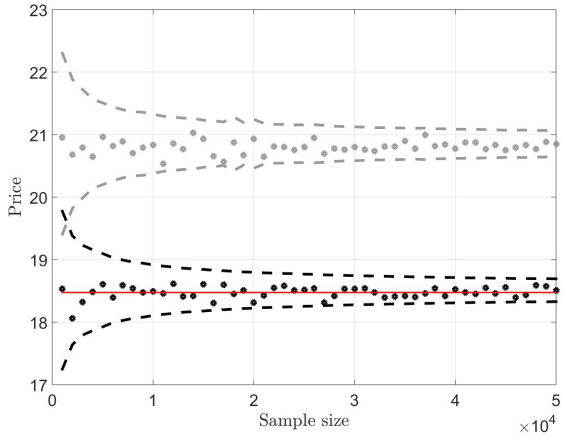

(c) Heston driven VTS used as an underlying with $\sigma_{target} = 12\%$

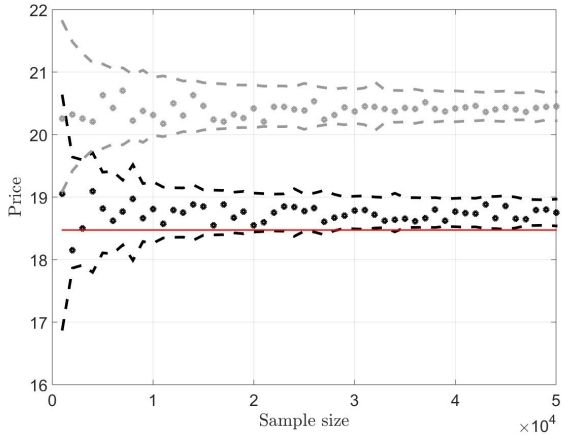

(d) Bates driven VTS used as an underlying with $\sigma_{target} = 12\%$

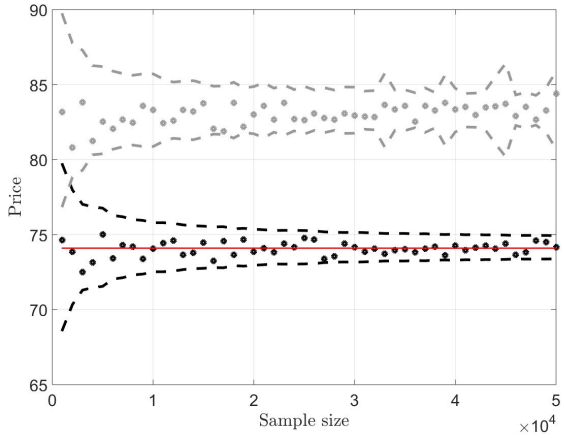

(e) Heston driven VTS used as an underlying with $\sigma_{target} = 30\%$

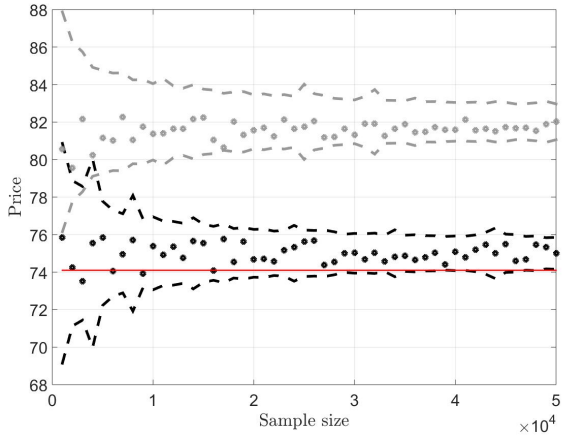

(f) Bates driven VTS used as an underlying with $\sigma_{target} = 30\%$

Fig. 4.1: Error corrected MCS prices for a floating strike lookback put where daily $(\Delta t = \delta t)$ and weekly $(\Delta t = 5\delta t)$ calibration has been used to construct the VTS, the standard deviation bounds and respective BS closed-form prices are also given

Chapter 5

Conclusions & Recommendations

This dissertation illustrates on a high level how stochastic volatility estimates can be used to adjust weightings of an actively managed portfolio to obtain underlying dynamics that make the process of pricing under BS more accurate. The result is also true for path-dependent options since the VTS is path-dependent itself. It was shown that the accuracy of the VTS is heavily influenced by the rate at which the portfolio weightings are adjusted, in other words how often model parameters are obtained from calibration, since these parameters are used to estimate the level of volatility inherent in the risky asset. It was found that daily calibration can obtain underlying dynamics that tend very close to GBM and thus an underlying that obeys the laws of BS. This holds true for both SVMs tested in this dissertation, being the Heston and Bates models. As the calibration rate is reduced (weekly, monthly or quarterly calibration) it was observed that the underlying dynamics become more like those of the risky asset that drive the portfolio. Nonetheless even a calibration rate as infrequent as quarterly calibration still yields somewhat improved implied volatilities as illustrated by Fig. 3.4. Another factor that was shown to heavily influence the accuracy of the VTS was the level of target volatility chosen. A high target volatility would usually exceed the measured risky asset volatility leading to an overall volatile portfolio that deviates from GBM. The converse true for low target volatilities. This is a direct consequence of the volatility of variance that exists in the risky asset.

It appears that under the assumptions made, the VTS manages to achieve its desired objectives, but it is important to understand the implications of relaxing these assumptions. Firstly it is a known fact that interest rates are not deterministic but stochastic themselves (to a lesser extent than assets but still stochastic nonetheless). The result means that a portion of volatility will be allocated to the VTS even when it selects its risk-free asset. The correlation that exists between the risky asset and the risk-free asset would need to be incorporated into the dynamics of the VTS, adding to the complexity of the problem. Secondly, as previously mentioned the risky asset of the VTS could typically comprise of a dividend paying ETF. Dividends would introduce another stochastic component into the VTS dynamics which would further complicate the problem. Finally, even though it was shown that daily calibration yields the desired return properties, it does so at the expense of a high calibration rate, which would lead to high transactional costs introducing drag into the portfolio. These problem areas would need to be addressed in future research before the VTS could be implemented realistically.

Bibliography

Anderson, L. and Brotherton-Ratcliffe, R. (1996). Exact exotics, *Risk* (9): 85–89.

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

Chapra, S. C. and Canale, R. P. (1998). *Numerical methods for engineers*, Vol. 2, Mcgraw-hill New York.

Conze, A. (1991). Path dependent options: The case of lookback options, *The Journal of Finance* **46**(5): 1893–1907.

Dumas, B., Fleming, J. and Whaley, R. E. (1998). Implied volatility functions: Empirical tests, *The Journal of Finance* **53**(6): 2059–2106.

Fama, E. F. (1998). Market efficiency, long-term returns, and behavioral finance, *Journal of financial economics* **49**(3): 283–306.

Ge, W. (2016). A survey of three derivative-based methods to harvest the volatility premium in equity markets, *The Journal of Investing* **25**(3): 48–58.

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

Hocquard, A., Ng, S. and Papageorgiou, N. (2013). A constant-volatility framework for managing tail risk, *Journal of Portfolio Management* **39**(2): 28.

Homescu, C. (2011). Implied volatility surface: construction methodologies and characteristics, *SSRN Electronic Journal.10.2139/ssrn.1882567* .

Hull, J. *et al. (2009). Options, futures and other derivatives/john c. hull.*

Israelov, R. and Klein, M. (2016). Risk and return of equity index collar strategies, The Journal of Alternative Investments **19***(1): 41–54.*

Kienitz, J. and Wetterau, D. (2012). Financial modelling: Theory, implementation and practice with MATLAB source*, John Wiley & Sons.*

Mandelbrot, B. B. (1997). The variation of certain speculative prices, pp. 371–418.

Papageorgiou, N. A., Reeves, J. J. and Sherris, M. (2017). Equity investing with targeted constant volatility exposure.

Rouah, F. D. (2013). The Heston Model and Its Extensions in Matlab and C#, *Wiley.*

Shreve, S. E. (2004). Stochastic calculus for finance II: Continuous-time models, *Vol. 11, Springer Science & Business Media.*

Stutzer, M. (1996). A simple nonparametric approach to derivative security valuation, The Journal of Finance **51***(5): 1633–1652.*

Appendix A

Proof of Equation (4.3)

Let $X_t = \alpha + \mu t + \sigma W_t$, be an arithmetic Brownian motion process over time interval $0 \leq s \leq t$. We are trying to prove that the probability of the running maximum $X_t^* = \sup X_s$ being at least greater than or equal to the absolute maximum y, conditional on the change of X_t over interval $\mathrm{d}x$ is equal to

$$\mathbb{P}\left(X_t^* \geq y | X_t \in \mathrm{d}x\right) = \exp\left(\frac{-2(y-x)(y-\alpha)}{\sigma^2 t}\right) \tag{A.1}$$

Proof. Equation (A.1) can be redefined, by using the laws of conditional probability, as

$$\begin{aligned}\mathbb{P}\left(X_t^* \geq y | X_t \in \mathrm{d}x\right) &= \frac{\mathbb{P}\left(X_t^* \geq y, X_t \in \mathrm{d}x\right)}{\mathbb{P}\left(X_t \in \mathrm{d}x\right)} \\ &= \frac{\frac{\partial}{\partial x}\mathbb{P}\left(X_t^* \geq y, X_t \leq x\right)}{\frac{\partial}{\partial x}\mathbb{P}\left(X_t \leq x\right)}\end{aligned} \tag{A.2}$$

The expression $\mathbb{P}\left(X_t \leq x\right)$ in the denominator of equation (A.2) can simply be defined as

$$\mathbb{P}\left(X_t \leq x\right) = \mathcal{N}\left(\frac{x - (\alpha + \mu t)}{\sigma\sqrt{t}}\right) \tag{A.3}$$

where $\mathcal{N}()$ is the normal cumulative distribution function.

To determine the expression $\mathbb{P}\left(X_t^* \geq y, X_t \leq x\right)$ that is found in the numerator of (A.2) we now define $\widetilde{\mathbb{P}}$ to be a measure obtained from $\mathbb{P}$ using Girsanov kernel $-\frac{\mu}{\sigma}$. Then

$$\widetilde{W_s} = W_s + \frac{\mu}{\sigma}s$$

is a $\widetilde{\mathbb{P}}$-Brownian motion. Our process now becomes

$$X_s - \alpha = \mu s + \sigma\left(\widetilde{W_s} - \frac{\mu}{\sigma}s\right) = \sigma\widetilde{W_s}$$

Therefore $\frac{X_s - \alpha}{\sigma}$ is a standard $\widetilde{\mathbb{P}}$-Brownian motion for $0 \leq s \leq t$. To revert back from $\widetilde{\mathbb{P}}$ to $\mathbb{P}$ we must add $\frac{\mu}{\sigma}$ to the drift. This is accomplished with a Girsanov transformation where the kernel is equal to $\frac{\mu}{\sigma}$, the result is

$$\frac{\mathrm{d}\widetilde{\mathbb{P}}}{\mathrm{d}\mathbb{P}} = \exp\left(-\frac{\mu}{\sigma}W_t - \frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2 t\right)$$

or

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\widetilde{\mathbb{P}}} = \exp\left(+\frac{\mu}{\sigma}\widetilde{W}_t - \frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2 t\right) \tag{A.4}$$

Initially it is assumed that $\sigma = 1$ and $\alpha = 0$. Then $X_s = \widetilde{W}_s$ is a $\widetilde{\mathbb{P}}$-Brownian motion for $0 \leq s \leq t$ and $\frac{\mu}{\sigma}$ simplifies to μ. Next we define an event

$$A = \{\omega \in \Omega : \widetilde{W_t^*}(\omega) \geq y, \widetilde{W}_t(\omega) \leq x\}$$

where we wish to obtain $\mathbb{P}(A)$.

Let $\widetilde{Z}_t$ be the $\widetilde{\mathbb{P}}$-Brownian motion obtained from $\widetilde{W}_t$ by reflecting it once it hits y. Next we define an event

$$B = \{\omega \in \Omega : \widetilde{Z_t^*}(\omega) \geq y, \widetilde{Z_t}(\omega) \leq x\}$$

Shreve (2004) shows by the reflection theorem that the above event can be redefined as

$$B = \{\omega \in \Omega : \widetilde{W}_t(\omega) \geq 2y - x\} \tag{A.5}$$

It can be assumed that

$$\widetilde{\mathbb{E}}\left[\mathbb{I}_{\{A\}} \exp\left(\mu\widetilde{W}_t - \frac{1}{2}\mu^2 t\right)\right] = \widetilde{\mathbb{E}}\left[\mathbb{I}_{\{B\}} \exp\left(\mu\widetilde{Z}_t - \frac{1}{2}\mu^2 t\right)\right] \tag{A.6}$$

because $\widetilde{Z}_t$ and $\widetilde{W}_t$ are identically distributed under $\widetilde{\mathbb{P}}$. It follows that using equation (A.4) and relations (A.5) and (A.6) yields

$$\begin{aligned}
\mathbb{P}(A) &= \widetilde{\mathbb{E}}\left[\mathbb{I}_{\{A\}} \frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\widetilde{\mathbb{P}}}\right] \\
&= \widetilde{\mathbb{E}}\left[\mathbb{I}_{\{A\}} \exp\left(\mu\widetilde{W}_t - \frac{1}{2}\mu^2 t\right)\right] \\
&= \widetilde{\mathbb{E}}\left[\mathbb{I}_{\{B\}} \exp\left(\mu\widetilde{Z}_t - \frac{1}{2}\mu^2 t\right)\right] \\
&= \widetilde{\mathbb{E}}\left[\mathbb{I}_{\{\widetilde{W}_t \geq 2y-x\}} \exp\left(\mu\left(2y - \widetilde{W}_t\right) - \frac{1}{2}\mu^2 t\right)\right] \\
&= \exp[2\mu y]\, \widetilde{\mathbb{E}}\left[\mathbb{I}_{\{\widetilde{W}_t \geq 2y-x\}} \exp\left(-\mu\widetilde{W}_t - \frac{1}{2}\mu^2 t\right)\right]
\end{aligned}$$

Now the factor $\exp\left(-\mu\widetilde{W}_t - \frac{1}{2}\mu^2 t\right)$, in the last expectation, is the Radon-Nikodym derivative of measure $\widehat{\mathbb{P}}$ obtained from $\widetilde{\mathbb{P}}$ by a Girsanov transformation that changes the drift of $\widetilde{W}_s$ to $-\mu$, therefore

$$\widetilde{\mathbb{E}}\left[\mathbb{I}_{\{\widetilde{W}_t \geq 2y-x\}} \exp\left(-\mu\widetilde{W}_t - \frac{1}{2}\mu^2 t\right)\right] = \widehat{\mathbb{E}}\left[\mathbb{I}_{\{\tilde{W}_t \geq 2y-x\}}\right] = \widehat{\mathbb{P}}\left(\widetilde{W}_t \geq 2y - x\right)$$

but under $\widehat{\mathbb{P}}$, $\widetilde{W}_s$ has drift -μ therefore $\widehat{W}_s = \widetilde{W}_s + \mu s$ is a standard Brownian motion under $\widehat{\mathbb{P}}$.

It follows that

$$\widehat{\mathbb{P}}\left(\widetilde{W_t} \geq 2y - x\right) = \widehat{\mathbb{P}}\left(\widehat{W_t} \geq 2y - x + \mu t\right) = \mathcal{N}\left(\frac{x - 2y - \mu t}{\sqrt{t}}\right)$$

and that

$$\mathbb{P}(A) = \exp(2\mu y)\,\mathcal{N}\left(\frac{x - 2y - \mu t}{\sqrt{t}}\right)$$

Now if $\sigma \neq 1$ then $\frac{X_t}{\sigma}$ is, under measure $\mathbb{P}$, an arithmetic Brownian motion with drift $\frac{\mu}{\sigma}$ and variance 1. Therefore

$$\begin{aligned}\mathbb{P}(X_t^* \geq y, X_t \leq x) &= \mathbb{P}\left(\frac{X_t}{\sigma} \leq \frac{x}{\sigma}, \frac{X_t^*}{\sigma} \geq \frac{y}{\sigma}\right) \\ &= \exp\left(\frac{2\mu y}{\sigma^2}\right)\mathcal{N}\left(\frac{x - 2y - \mu t}{\sigma\sqrt{t}}\right)\end{aligned}$$

Next if $\alpha \neq 0$ then $X_s - \alpha$ is an arithmetic Brownian motion starting at 0, implying

$$\begin{aligned}\mathbb{P}(X_t^* \geq y, X_t \leq x) &= \mathbb{P}((X - \alpha)_t^* \geq y - \alpha, X_t - \alpha \leq x - \alpha) \\ &= \exp\left(\frac{2\mu(y - \alpha)}{\sigma^2}\right)\mathcal{N}\left(\frac{x - (2y - \alpha + \mu t)}{\sigma\sqrt{t}}\right) \qquad \text{(A.7)}\end{aligned}$$

Substituting equations (A.3) and (A.7) into equation (A.2) while simultaneously calculating the derivatives with respect to x and solving, yields

$$\begin{aligned}\mathbb{P}(X_t^* \geq y | X_t \in \mathrm{d}x) &= \frac{\left(\frac{1}{\sigma\sqrt{t}}\right)\exp\left(\frac{2\mu(y-\alpha)}{\sigma^2}\right)\varphi\left(\frac{x-(2y-\alpha+\mu t)}{\sigma\sqrt{t}}\right)}{\left(\frac{1}{\sigma\sqrt{t}}\right)\varphi\left(\frac{x-(\alpha+\mu t)}{\sigma\sqrt{t}}\right)} \\ &= \exp\left(\frac{-2(y-x)(y-\alpha)}{\sigma^2 t}\right)\end{aligned}$$

where $\varphi\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{\sqrt{2\pi\sigma^2}}\exp\left[\frac{-(x-\mu)^2}{2\sigma^2}\right]$ is the normal density distribution function. □