Volume 39(1), pp. 65 – 91
http://orion.journals.ac.za

ORiON
ISSN 0259–191X (print)
ISSN 2224–0004 (online)

On the calibration of stochastic volatility models to estimate the real-world measure used in option pricing

Alexis Levendis[*] and Eben Maré[†]

Received: 6 September 2022; Revised: 1 June 2023; Accepted: 7 June 2023

Abstract

It is widely noted that the Heston stochastic volatility model fails to capture the fat tails often observed in daily equity returns. Adding random jumps improves the model's ability to capture extreme events. This extension is known as the Bates stochastic volatility jump (SVJ) model. The model parameters for the Heston and Bates SVJ models are generally calibrated to option prices inducing the so-called risk-neutral measure. However, in the absence of a sufficiently liquid options market, one has to resort to calibration under the real-world measure. In this paper, we calibrate the Heston and Bates SVJ models to historical equity returns in the United States and South Africa using the efficient method of moments (EMM). We then show how a real-world stochastic volatility model can be used in practice to test a simple volatility targeting strategy. Our findings suggest that stochastic volatility and jumps are both required to characterise equity returns in South Africa. Furthermore, volatility targeting is an effective strategy that allows investors to manage the downside risk of a portfolio.

Key words: Heston model, Bates stochastic volatility jump model, calibration, real-world measure, efficient method of moments, volatility targeting

1 Introduction

Equity prices tend to fall sharply in times of economic crisis. Think back to 19 October 1987, *Black Monday*, when the Dow plummeted more than 20% in a single day, for example.

[*]Department of Actuarial Science, University of Pretoria, Private bag X20, Hatfield, 0028, Republic of South Africa, fax: 021 8082409, email: alexilevendis@gmail.com
[†]Department of Mathematics and Applied Mathematics, University of Pretoria

http://dx.doi.org/10.5784/39-1-747

Whether it be war, a global pandemic, or a great recession, there is no way of knowing when a market crash will occur.

The seminal Black-Scholes [6] model used for the pricing of contingent claims, for example, is based on the assumption of log-normal asset returns – the driving force of asset returns in this model assumes a geometric Brownian motion (gBm). Such models are flawed in this respect since they do not account for severe market shocks or the stochastic nature of asset return volatility.

Strides have been made in the asset-pricing literature to account for the non-normality of equity returns (see, Cont [8], for a discussion on statistical properties of asset returns). The Heston [14] stochastic volatility model, for example, is frequently used in the pricing of contingent claims based on equities. It provides a reasonably good calibrated fit to the long-term implied volatility skew observed in the equity derivatives market (see, [12]). However, for short maturities, the model fails to produce the steep slope of the skew (see, for example, [9]). This indicates that the model is unable to account for short-term shocks (so-called jumps) that may be caused by adverse market events. Statistical tests have been proposed in [1] and [2] to verify whether jumps are present in financial time series.

A well-established model that can provide a good fit to the short-term skew is the Bates [5] stochastic volatility jump (SVJ) model. The Bates SVJ model is an extension of the classical Heston [14] model that adds random jumps based on a Poisson process. Poklewski-Koziell [17] performed a detailed analysis of the Heston and Bates SVJ models and showed that the Bates SVJ model produces a good fit to the S&P500 implied volatility surface when compared to the Heston model. Note that Poklewski-Koziell calibrated the Heston and Bates SVJ models to option prices (this calibration is said to be in the risk-neutral pricing measure).

Models with jump dynamics are better at characterising markets with significant implied volatility skews than models without jumps (see, [12]). However, past literature on jump diffusion models tends to focus exclusively on calibration to option prices (*i.e.*, the risk-neutral measure). In the absence of a liquidly traded options market, calibration using standard least-squares techniques to minimise the sum of squared differences between market and model prices is infeasible. This is a common problem that plagues illiquid option markets like South Africa. Another challenge is that the estimated density resulting from the calibration to option prices can differ substantially from the estimated density of the historical log returns (see, Visagie and Grobler [19]).

Calibration of continuous-time models to discrete historical observations has been explored by Andersen, Benzoni, and Lund [3]; they applied the efficient method of moments (EMM) technique of Gallant and Tauchan [11] to calibrate a class of SVJ models with Poisson jumps of time-varying intensity to daily S&P500 returns. The EMM is a simulation-based technique that estimates the continuous-time model parameters from the expectation of the derivative of the log-likelihood function where the log-likelihood function takes the form of a simpler discrete-time auxiliary semi-non-parametric model. Andersen *et al.* [3] conclude that stochastic volatility and jumps are both important factors required to characterise daily S&P500 returns.

To our knowledge, there has been no attempt made to calibrate stochastic volatility models, including the Heston and Bates SVJ models, to historical equity returns in South Africa (*i.e.*, calibration to the so-called real-world measure). The literature tends to focus solely on calibration to option prices (risk-neutral measure), which, in the context of South African single stocks, for example, is typically impractical (as no generally liquid traded market exists). We, therefore, view this as an excellent opportunity to contribute to the asset-pricing literature with specific focus on calibration of stochastic volatility models in illiquid markets.

The real-world measure is often neglected in favour of the risk-neutral measure due to the pricing of contingent claims. However, the real-world measure is extremely useful and important in risk-management and asset/liability applications (see, *e.g.*, [18]); simulation-based analysis of trading and investment strategies (see, [16]); analysis of so-called xVA (counterparty credit, margin, and capital costs); investment based pricing and evaluation of asset price behaviour; and product development, to name a few.

The goal of this paper is to calibrate the Heston and Bates SVJ models to historical FTSE/JSE Top40 returns by making use of the EMM and to test which model is better at characterising the evolution of equity-based risk and return in the South African market. Furthermore, we show how a real-world stochastic volatility model can be used in the context of portfolio management by testing a simple volatility targeting strategy.

The remainder of this paper is structured as follows, Section 2 presents the dynamics for the Heston and Bates SVJ models as well as the EMM technique of Gallant and Tauchan. Section 3 shows the calibration results of the EMM applied to the S&P500 and FTSE/JSE Top40. Section 4 focuses on a practical application of real-world stochastic volatility models by testing a volatility targeting strategy, and Section 5 concludes the paper.

2 Stochastic volatility models

This section is divided into three parts. Subsection 1 is dedicated to the Heston stochastic volatility model. Subsection 2 focuses on the Bates SVJ model, and Subsection 3 discusses the EMM methodology of Gallant and Tauchan.

2.1 The Heston stochastic volatility model

Under the real-world probability measure, $\mathbb{P}$, the Heston model is given by the system of stochastic differential equations (SDEs)

$$
\begin{aligned}
dS(t) &= \mu S(t)dt + \sqrt{v(t)}S(t)dW_x(t), \\
dv(t) &= \big(\alpha - \beta v(t)\big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \\
dW_x(t)dW_v(t) &= \rho_{x,v}dt, \qquad (1)
\end{aligned}
$$

where μ denotes the expected rate of return, β is the mean reversion speed of the variance, $\frac{\alpha}{\beta}$ is the long-run mean of the variance, σ_v is the volatility of the variance, and $\rho_{x,v}$ is the correlation between the stock and variance processes.

Taking the log-transform $x(t) = \log S(t)$ and applying Itô's lemma to the equations in (1), we obtain

$$\begin{aligned} dx(t) &= \big(\mu - \frac{1}{2}v(t)\big)dt + \sqrt{v(t)}dW_x(t), \\ dv(t) &= \big(\alpha - \beta v(t)\big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt. \end{aligned} \tag{2}$$

We simulate 100,000 paths over 100 daily time steps from (2) using an Euler discretisation scheme to illustrate the impact of certain Heston model parameters on the returns distribution at the end of 100 days. Figure 1 shows the impact of σ_v.

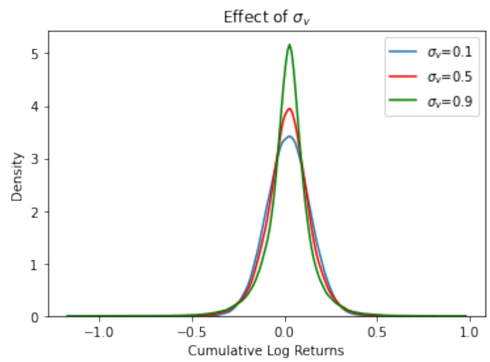

Figure 1: *Effect of σ_v on returns.*

Note that the kurtosis of the returns distribution increases as σ_v increases (kurtosis values of 3.13, 5.90, and 12.18 corresponding to σ_v values of 0.1, 0.5, and 0.9). Figure 2 illustrates the impact of $\rho_{x,v}$.

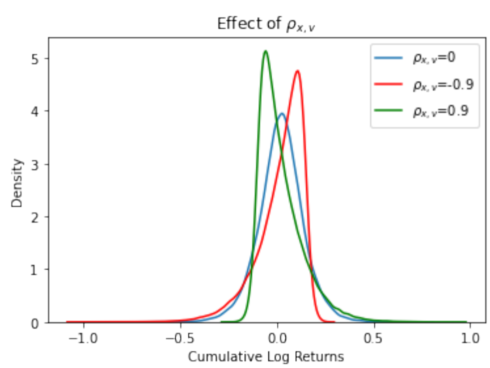

Figure 2: *Effect of $\rho_{x,v}$ on returns.*

The parameter $\rho_{x,v}$ controls the skewness of the returns distribution – a negative value for $\rho_{x,v}$ produces a negatively skewed distribution and vice versa.

Even though the Heston model displays some desirable properties, it still underestimates the kurtosis of returns observed in practice (see, [13]). Therefore, we consider the Bates SVJ model next.

2.2 The Bates stochastic volatility jump model

Under the $\mathbb{P}$-measure, the Bates SVJ model is presented by the system of SDEs

$$\begin{aligned} dS(t) &= \big(\mu - \lambda\mu_S\big)S(t)dt + \sqrt{v(t)}S(t)dW_x(t) + JS(t)dN(t), \\ dv(t) &= \big(\alpha - \beta v(t)\big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt, \end{aligned} \tag{3}$$

where $N(t)$ is a Poisson process with jump intensity λ. Furthermore, J denotes the percentage jump size of the underlying where

$$\log(1+J) \sim \mathcal{N}\Big(\log(1+\mu_S) - 0.5\sigma_S^2, \sigma_S^2\Big),$$

with μ_S and σ_S the mean and volatility of the jump size, and $\mathcal{N}(\cdot,\cdot)$ a normally distributed random variable.

Taking the log-transform $x(t) = \log S(t)$ and applying Itô's lemma to the equations in (3), we get

$$\begin{aligned} dx(t) &= \big(\mu - \lambda\mu_S - \frac{1}{2}v(t)\big)dt + \sqrt{v(t)}dW_x(t) + \log(1+J)dN(t), \\ dv(t) &= \big(\alpha - \beta v(t)\big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt. \end{aligned} \tag{4}$$

Similar to what was done using the Heston model, we simulate 100,000 paths over 100 daily time steps from the equations in (4) to illustrate the impact of the jump parameters on the returns distribution. Figure 3 shows the impact of λ.

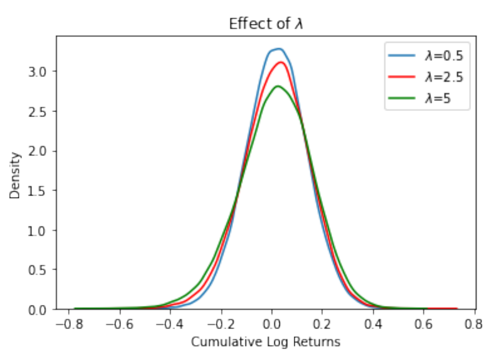

Figure 3: *Effect of λ on returns.*

Note that the kurtosis of the distribution increases as λ increases (kurtosis values of 3.12, 3.14, and 3.18 corresponding to λ values of 0.5, 2.5, and 5). Figure 4 illustrates the impact of the mean jump size μ_S.

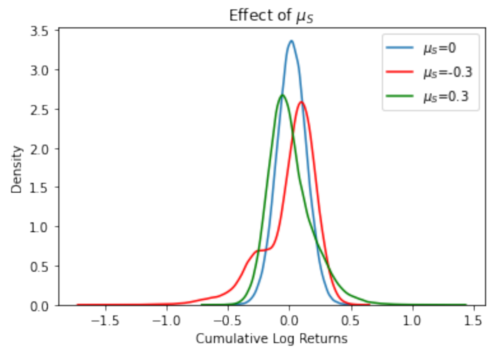

Figure 4: *Effect of μ_S on returns.*

Similar to the parameter $\rho_{x,v}$, the mean jump size μ_S controls the skewness of the distribution – a negative value for μ_S produces a negatively skewed distribution and vice versa. Lastly, the Figure 5 shows the impact of the volatility of the jump size σ_S.

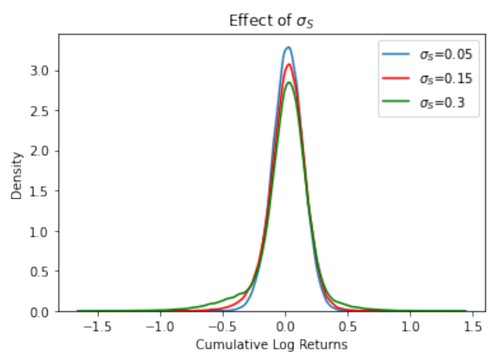

Figure 5: *Effect of σ_S on returns.*

Note that the kurtosis of the distribution increases as σ_S increases (kurtosis values of 6.08, 6.21, and 6.34 corresponding to σ_S values of 0.05, 0.15, and 0.3).

It should be clear that the Bates SVJ model provides more flexibility than the Heston model to capture additional features of returns. In the next subsection, we present the EMM technique of Gallant and Tauchan that will be used to calibrate the Heston and Bates SVJ models to historical equity returns.

2.3 Efficient method of moments

The EMM procedure is outlined in Andersen, Chung, and Sorensen [4]. Suppose there is a historical time series $Y_t = \{y_1, ...y_t\}$ and the aim is to estimate the vector of stochastic volatility model parameters $\boldsymbol{\theta}$ from this series. There is no analytical expression for the likelihood function of the Heston or Bates SVJ models. Therefore, the first step in the EMM procedure is to choose an auxiliary model (score generator) with transition density function $f(y_t|Y_{t-1}, \boldsymbol{\eta})$ parameterised by the pseudo parameter vector $\boldsymbol{\eta}$. To this end, we choose the semi-nonparametric density of Gallant and Nychka [10] where a leading parametric term is chosen to capture the majority of the dependency in the conditional mean and variance. The transition density function is then extended by adding Hermite polynomials that capture any remaining non-Gaussian features in the time series. The semi-nonparametric density is given by

$$f_K(y_t|Y_{t-1}; \boldsymbol{\eta}) = \left(v + (1-v) \frac{[P_K(z_t)]^2}{\int_{\mathbb{R}} [P_K(u)]^2 \phi(u) du} \right) \frac{\phi(z_t)}{\sqrt{h_t}},$$

where $v = 0.01$ to avoid instability during the EMM estimation when $P_K(z_t) = 0$, $\phi(\cdot)$ is the standard normal density function, and

$$z_t = \frac{y_t - \mu_t}{\sqrt{h_t}},$$

with

$$\mu_t = 0,$$

$$h_t = \omega + \gamma_0 y_{t-1}^2 + \gamma_1 h_{t-1} \sim GARCH(1,1),$$

where μ_t and h_t denote the conditional mean and variance, respectively. Note that μ_t need not be zero, and h_t may be specified by alternative discrete-time models like ARCH or EGARCH (see, [4], for details).

The Hermite polynomials are given by:

$$P_K(z_t) = \sum_{i=0}^{K_z} a_i z_t^i,$$

where $a_0 = 1$ and K_z denotes the order of the Hermite polynomial. Note that the EMM procedure requires the dimension of the pseudo parameter vector $\boldsymbol{\eta}$ to be greater than or equal to the dimension of the stochastic volatility parameter vector $\boldsymbol{\theta}$.

Once an auxiliary model has been chosen, the pseudo parameter vector $\boldsymbol{\eta}$ is estimated using maximum likelihood. The maximum likelihood estimator $\hat{\boldsymbol{\eta}}_T$ satisfies the first-order conditions

$$\frac{1}{T} \sum_{t=1}^{T} \frac{\partial}{\partial \boldsymbol{\eta}} \log f(y_t|Y_{t-1}, \hat{\boldsymbol{\eta}}_T) = \frac{1}{T} \sum_{t=1}^{T} s_f(Y_t, \hat{\boldsymbol{\eta}}_T) = 0,$$

where $s_f(Y_t, \hat{\boldsymbol{\eta}}_T) = \frac{\partial}{\partial \boldsymbol{\eta}} \log f(y_t|Y_{t-1}, \hat{\boldsymbol{\eta}}_T)$ denotes the score function of the auxiliary model.

The second step in the EMM procedure is to simulate a series $\hat{y}_n(\boldsymbol{\theta})$, $n = 1, ..., N$, from the stochastic volatility model for a given $\boldsymbol{\theta}$ and evaluate the sample moments at the fixed maximum likelihood estimate $\hat{\boldsymbol{\eta}}_T$, *i.e.*,

$$m_N(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}_T) = \frac{1}{N} \sum_{n=1}^{N} \frac{\partial}{\partial \boldsymbol{\eta}} \log f(\hat{y}_n(\boldsymbol{\theta})|\hat{Y}_{n-1}(\boldsymbol{\theta}), \hat{\boldsymbol{\eta}}_T).$$

Andersen *et al.* [4] mention that $m_N(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}_T) \to m(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}_T)$ as $N \to \infty$; hence, the simulated time series should be large enough so that the Monte Carlo error can be ignored.

Finally, the vector of stochastic volatility parameters $\boldsymbol{\theta}$ is estimated by minimising the objective function

$$\varphi = \underset{\boldsymbol{\theta}}{\operatorname{argmin}}[m_N(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}_T)' \hat{I}_T^{-1} m_N(\boldsymbol{\theta}, \hat{\boldsymbol{\eta}}_T)], \tag{5}$$

where $\hat{I}_T$ is a consistent estimator of the asymptotic covariance matrix I of the sample pseudo score vector. The estimator $\hat{I}_T$ is calculated as the outer product of scores, *i.e.*,

$$\hat{I}_T = \frac{1}{T}\sum_{t=1}^{T} \frac{\partial}{\partial \boldsymbol{\eta}} \log f(y_t|Y_{t-1}, \hat{\boldsymbol{\eta}}_T) \frac{\partial}{\partial \boldsymbol{\eta}} \log f(y_t|Y_{t-1}, \hat{\boldsymbol{\eta}}_T)'.$$

A major advantage of the EMM is that T multiplied by the minimised value in (5) follows a χ^2 distribution with $n_{\boldsymbol{\eta}} - n_{\boldsymbol{\theta}}$ degrees of freedom where $n_{\boldsymbol{\eta}}$ and $n_{\boldsymbol{\theta}}$ denote the number of parameters in the semi-nonparametric and stochastic volatility models. Therefore, a goodness-of-fit test can be performed by comparing the final estimate of the objective function multiplied by the number of observations used in the calibration with the relevant percentile from the χ^2 distribution. Andersen *et al.* [4] explain that the null hypothesis in this case is that the model has been correctly specified. Hence, if the minimised value in (5) multiplied by the number of observations is less than the critical value from the χ^2 distribution, the hypothesis is not rejected.

In the next section, we calibrate the Heston and Bates SVJ models to historical S&P500 and FTSE/JSE Top40 returns using the EMM.

3 Empirical results

This section is divided into two parts. The first subsection focuses on the S&P500 and compares our EMM implementation of the Heston and Bates SVJ models to the results in Andersen *et al.* [3]. The reason for this is to validate the accuracy of our implementation. We further extend the analysis of Andersen *et al.* [3] by calibrating the Heston model over different periods to test the stability of the model parameters. The second subsection uses the EMM to calibrate the Heston and Bates SVJ models to the FTSE/JSE Top40 to test which model is better at capturing risk and return in the South African equity market.

3.1 S&P500

The S&P500 is considered by many to be the best indicator of global equity market performance. It consists of 500 leading companies that are publicly traded in the United States. Andersen *et al.* [3] calibrated the Heston and Bates SVJ models using the EMM over the period 02/01/1953 to 31/12/1996 to S&P500 returns. Figure 6 below shows the daily historical closing prices for the S&P500 from 02/01/1953 to 31/12/1996.

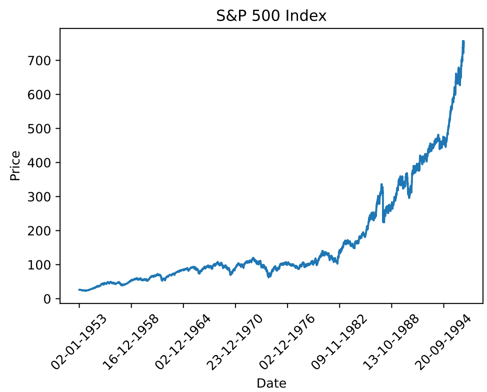

Figure 6: *S&P500 historical closing prices.*

Figure 7 below shows the daily log returns for the S&P500 from 02/01/1953 to 31/12/1996.

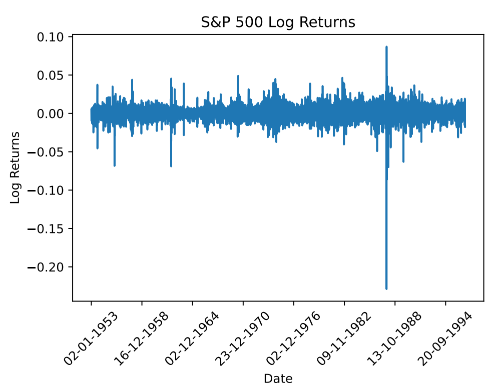

Figure 7: *S&P500 log returns.*

The most significant event over the period 02/01/1953 to 31/12/1996 was *Black Monday*, 19 October 1987, when the S&P500 fell by more than 20%. Prior to *Black Monday*, equity market behaviour was relatively stable. The concept of a *volatility skew* was unknown and the assumption of log-normal returns seemed reasonable. Figure 8 below shows the daily S&P500 log returns with the normal distribution superimposed over the returns.

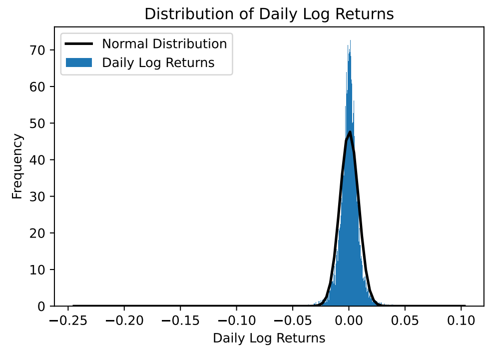

Figure 8: *S&P500 and normal density.*

It is clear from Figure 8 that the normal distribution cannot capture the high peaks and fat tails of the S&P500 returns (sample kurtosis of 60.08). Assuming a log-normal distribution of returns is, therefore, not consistent with historical equity behaviour and significantly underestimates the size and frequency of equity price drops.

Our first goal is to reproduce the calibration results in Andersen *et al.* [3] for the Heston and Bates SVJ models over the same period from 02/01/1953 to 31/12/1996 to validate the accuracy of our EMM implementation. The results for the Heston stochastic volatility model are shown below.

Heston stochastic volatility model

Table 1 below shows a comparison of the Heston model parameters based on our EMM implementation to the parameters in Andersen *et al.* [3].

Parameter	Andersen *et al.*	Our implementation
μ	0.0756	0.0768
α	0.0438	0.0485
β	3.2508	4.2142
σ_v	0.1850	0.1768
$\rho_{x,v}$	-0.5877	-0.4323
$T\varphi$	31.9400	15.6099

Table 1: *Comparison of annualised calibrated Heston parameters.*

The calibrated parameters achieved using our method and those of Andersen *et al.* [3] align well. We do not expect to match their parameters exactly since we are likely using a different optimisation routine and have access to better software compared to what was available 20 years ago. The goodness-of-fit statistic, $T\varphi$, is also shown in Table 1 for completeness.

Interestingly, the goodness-of-fit statistic is significantly smaller in our implementation than in Andersen *et al.* [3]. This indicates that our implementation of the Heston model yields a better fit to the S&P500 returns than the Heston model in Andersen *et al.* [3].

Figures 9 and 10 below show a visual comparison of the densities generated by the model parameters in Table 1.

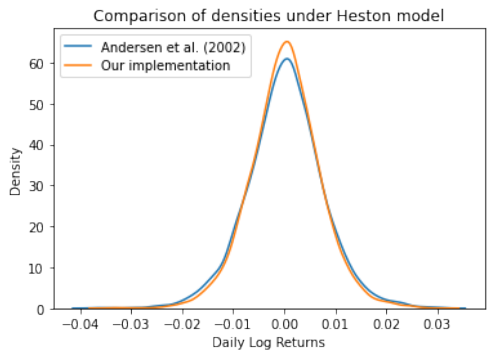

Figure 9: *Comparison of densities under the calibrated Heston stochastic volatility model.*

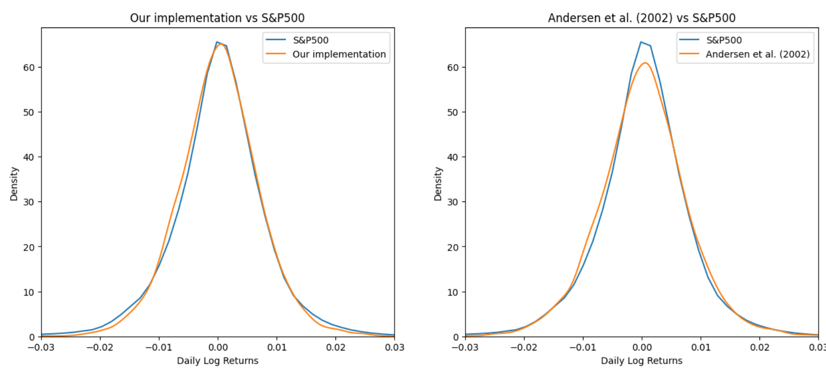

Figure 10: *Comparison of calibrated Heston models and S&P500 densities.*

Figure 9 indicates that the two sets of model parameters yield very similar results. Figure 10 shows that our implementation of the Heston model captures the peak of the S&P500 returns slightly better than Andersen *et al.* [3].

Table 2 below shows the first four statistical moments and the minimum/maximum values generated by the Heston model compared to the S&P500.

Statistic	S&P500 index	Andersen *et al.*	Our implementation
Mean	0.0301%	0.0231%	0.0235%
Std dev	0.8346%	0.7498%	0.6880%
Skewness	-2.0220	-0.0179	-0.0066
Kurtosis	60.0830	4.0254	3.9219
Minimum	-0.2290	-0.0379	-0.0349
Maximum	0.0871	0.0319	0.0310

Table 2: *Calibrated Heston model daily statistics for the S&P500.*

The first two moments of the S&P500 (mean and standard deviation) are captured relatively well by the Heston model. However, the Heston model substantially underestimates the skewness and kurtosis over the period 02/01/1953 to 31/12/1996. The minimum and maximum values are calculated by simulating a single trajectory for the log returns from the Heston model over a long period. The minimum/maximum is then calculated from the simulated return series. As explained by Andersen *et al.* [3], stochastic volatility on its own is not adequate to describe the observed S&P500 returns.

Next, we extend the results of Andersen *et al.* [3] by calibrating the Heston model over different periods spanning approximately 20 years to test the stability of the model parameters. Table 3 below shows the first four statistical moments and the minimum and maximum values over each period.

Statistic	1950-1970	1960-1980	1970-1990	1980-2000	1990-2010	2000-2022
Mean	0.0342%	0.0117%	0.0264%	0.0520%	0.0224%	0.0174%
Std dev%	0.6740%	0.7508%	0.9861%	0.9996%	1.1724%	1.2474%
Skewness	-0.6828	0.04194	-2.5083	-2.6223	-0.1980	-0.4009
Kurtosis	11.6208	7.5054	64.9265	62.9018	12.1684	13.5220
Minimum	-0.0691	-0.0691	-0.2290	-0.2290	-0.0947	-0.1277
Maximum	0.0454	0.0490	0.0871	0.0871	0.1096	0.1096

Table 3: *S&P500 daily statistics over different periods.*

The periods 1970 to 1990 and 1980 to 2000 both contain *Black Monday* and show substantially different values for the skewness and kurtosis compared to the other periods. Not even the global financial crisis of 2008 or the COVID-19 pandemic came close to the crash of 19 October 1987. Table 4 below shows the calibrated parameters and goodness-of-fit statistic for the Heston model using the EMM for each period.

Parameter	1950-1970	1960-1980	1970-1990	1980-2000	1990-2010	2000-2022
μ	0.0816	0.0277	0.0670	0.1359	0.0891	0.0858
α	0.0744	0.0534	0.0769	0.0774	0.0602	0.1137
β	8.4560	4.8925	5.0140	5.0555	3.0968	4.9395
σ_v	0.2574	0.2341	0.2015	0.2266	0.2292	0.3747
$\rho_{x,v}$	-0.5416	-0.5703	0.0347	-0.4144	-0.8891	-0.8801
$T\varphi$	5.6453	2.9665	6.7840	22.8005	18.1143	7.2800
$\chi^2_{0.05}$	9.4877	9.4877	9.4877	9.4877	9.4877	9.4877

Table 4: *Calibrated Heston model parameters (annualised).*

The expected return, μ, can vary substantially between periods and follows a similar trend to the mean of the S&P500 returns in Table 3. The Heston stochastic volatility parameters (α, β, σ_v) are relatively stable over time. The correlation parameter, $\rho_{x,v}$, shows a strong negative relationship between returns and asset volatility except for the period 1970 to 1990. Interestingly, $\rho_{x,v}$ changes from 0.0347 in the period 1970 to 1990 to -0.4144 in 1980 to 2000. Both these periods include *Black Monday.* This instability suggests that the estimation of $\rho_{x,v}$ is sensitive to outliers.

The goodness-of-fit statistic shows interesting results. From Table 3, the goodness-of-fit statistic suggests that the hypothesis that the observed data are realised from the calibrated Heston model is not rejected for the periods 1950 to 1970, 1960 to 1980, 1970 to 1990, and 2000 to 2022. However, this hypothesis is rejected for the periods 1980 to 2000 and 1990 to 2010, which indicates that the calibration is not stable over time.

In the next section, we compare our implementation of the Bates SVJ model to the reslts of Andersen *et al.* [3].

Bates stochastic volatility jump model

The Bates SVJ model adds three additional jump parameters, λ, μ_S, and σ_S to the standard Heston model. Andersen *et al.* [3] explain that the mean jump parameter, μ_S, is of less importance and poorly identified in general. Therefore, we follow Andersen *et al.* [3] by imposing the restriction $\mu_S = 0$. The Bates SVJ model in (4) then becomes

$$\begin{aligned} dx(t) &= \big(\mu - \frac{1}{2}v(t)\big)dt + \sqrt{v(t)}dW_x(t) + \log(1+J)dN(t), \\ dv(t) &= \big(\alpha - \beta v(t)\big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt, \end{aligned} \tag{6}$$

where

$$\log(1+J) \sim \mathcal{N}\Big(-0.5\sigma_S^2, \sigma_S^2\Big).$$

Table 5 below compares the Bates SVJ parameters based on our implementation of the EMM to that of the results in Andersen *et al.* [3] over the period 02/01/1953 to 31/12/1996.

Parameter	Andersen et al.	Our implementation
μ	0.0766	0.0815
α	0.0438	0.0481
β	3.0240	3.6782
σ_v	0.1792	0.2133
$\rho_{x,v}$	-0.6220	-0.4926
λ	5.0904	4.0147
σ_J	0.0134	0.0151
$T\varphi$	14.9000	2.3648

Table 5: *Comparison of annualised calibrated Bates SVJ parameters.*

Once again, the two sets of parameters align well. Based on our implementation, the goodness-of-fit statistic decreases from 15.6099 under the Heston model to 2.3648 under the Bates SVJ model. This indicates that both stochastic volatility and jumps are important factors to consider when modelling S&P500 returns; this aligns with the findings reported in Andersen *et al.* [3].

Similar to the Heston calibration in Table 1, an interesting observation is that the goodness-of-fit statistic based on our implementation is significantly smaller than Andersen *et al.* [3], which suggests a better fit to the S&P500 returns. Figures 11 and 12 below show a visual comparison of the densities generated by the model parameters in Table 5.

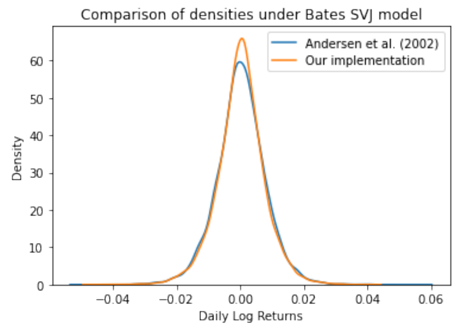

Figure 11: *Comparison of densities under the calibrated Bates SVJ model.*

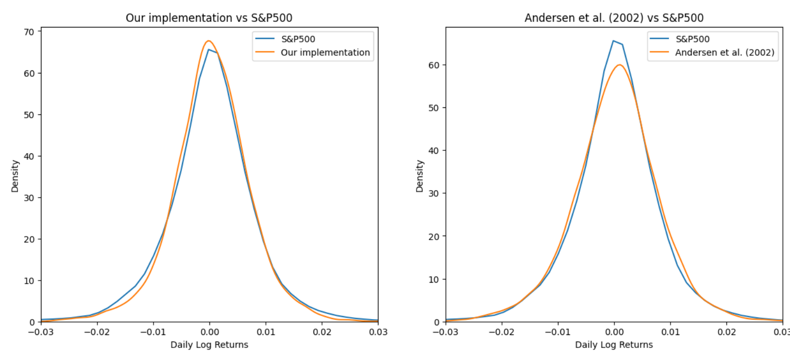

Figure 12: *Comparison of calibrated Bates models and S&P500 densities.*

Figure 11 shows that the distributions generated by the Bates SVJ model using our parameters versus the parameters in Andersen *et al.* [3] yield similar results. Figure 12 shows that our implementation captures the peak of the S&P500 returns slightly better than Andersen *et al.* [3]. Table 6 below shows the first four statistical moments and the minimum/maximum values generated by the Bates SVJ model compared with the S&P500.

Statistic	S&P500 index	Andersen *et al.*	Our implementation
Mean	0.0301%	0.0239%	0.0209%
Std dev	0.8346%	0.7717%	0.7433%
Skewness	-2.0220	-0.0133	-0.0923
Kurtosis	60.0830	4.6294	4.5946
Minimum	-0.2290	-0.0498	-0.0460
Maximum	0.0871	0.0566	0.0407

Table 6: *Calibrated Bates model daily statistics for the S&P500.*

The minimum and maximum values are calculated by simulating a single trajectory for the log returns from the Bates SVJ model over a long period. The minimum/maximum is then calculated from the simulated series. Adding jumps to the return process improves the results for the skewness and kurtosis, but not nearly enough to capture the severe market shock of 19 October 1987.

In the next section, we calibrate the Heston and Bates SVJ models to FTSE/JSE Top40 returns to test which model captures risk and return best in the South African market.

3.2 FTSE/JSE Top40

The FTSE/JSE Top40 is an index consisting of the 40 largest publicly traded companies by market capitalisation in South Africa. Figure 13 below shows daily historical closing prices for the FTSE/JSE Top40 from 30/06/1995 to 30/06/2022.

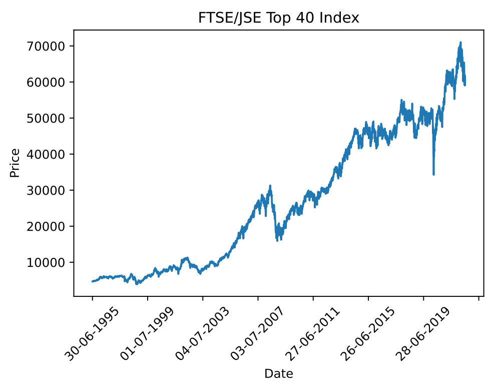

Figure 13: *FTSE/JSE Top40 historical closing prices.*

Figure 14 below shows the daily log returns for the FTSE/JSE Top40 from 30/06/1995 to 30/06/2022.

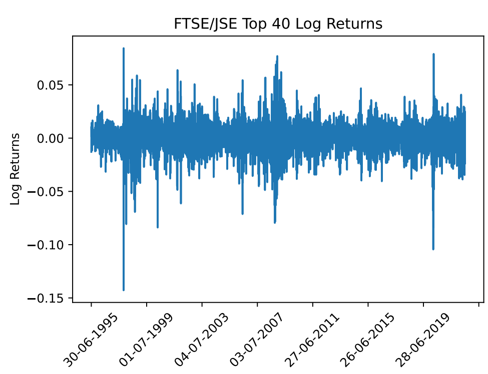

Figure 14: *FTSE/JSE Top40 log returns.*

Figure 15 below shows the normal distribution superimposed over the daily FTSE/JSE Top40 returns.

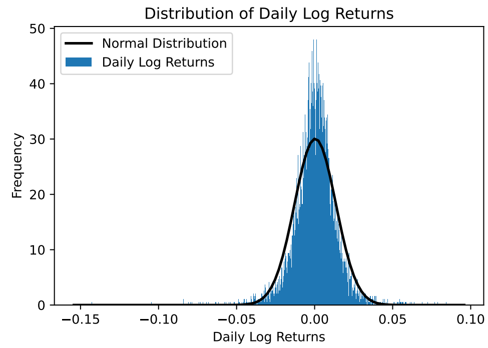

Figure 15: *FTSE/JSE Top40 and normal densities.*

Note that the normal distribution is not able to capture the high peak and fat tails observed in the empirical distribution of daily FTSE/JSE Top40 returns (sample kurtosis of 9.43). Next, we calibrate the Heston model to daily FTSE/JSE Top40 returns.

Heston stochastic volatility model

Table 7 below shows the calibrated Heston parameters to daily FTSE/JSE Top40 parameters over the period 30/06/1995 to 30/06/2022 using the EMM.

Parameter	Estimate
μ	0.0982
α	0.2342
β	6.9424
σ_v	0.4782
$\rho_{x,v}$	-0.9364
$T\varphi$	13.4381
$\chi^2_{0.05}$	9.4877

Table 7: *Annualised calibrated Heston parameters.*

Note that the asset-volatility correlation, $\rho_{x,v}$, for the FTSE/JSE Top40 is much more pronounced (negative) than the asset-volatility correlation for the S&P500. The goodness-of-fit statistic suggests that the hypothesis that the observed data are realised from the calibrated Heston model is rejected at a 5% level of significance.

Figure 16 compares the density generated by the Heston model parameters in Table 7 to a kernel density estimate of the FTSE/JSE Top40.

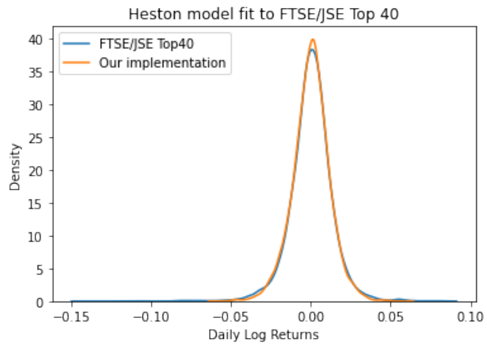

Figure 16: *Comparison of densities under the Heston model.*

Visually, the Heston model fits the FTSE/JSE Top40 density well. Table 8 below compares the first four moments and minimum/maximum values generated by the Heston model with the daily FTSE/JSE Top40 returns.

Statistic	FTSE/JSE Top40 index	Our implementation
Mean	0.0385%	0.0270%
Std dev	1.3290%	1.1679%
Skewness	-0.4369	-0.1313
Kurtosis	9.4344	4.1614
Minimum	-0.1429	-0.0585
Maximum	0.0845	0.0588

Table 8: *Calibrated Heston model daily statistics for the FTSE/JSE Top40.*

Similar to the results reported for the S&P500 in Table 2, the Heston model captures the mean and standard deviation of the FTSE/JSE Top40 well. However, the model underestimates the observed skewness and kurtosis. Next, we implement the Bates SVJ model to test if jumps improve the calibration.

Bates stochastic volatility jump model

Table 9 below shows the calibrated Bates SVJ parameters to daily FTSE/JSE Top40 returns over the period 30/06/1995 to 30/06/2022 using the EMM.

Parameter	Estimate
μ	0.0893
α	0.1815
β	5.4469
σ_v	0.4328
$\rho_{x,v}$	-0.8705
λ	3.9973
σ_J	0.0147
$T\varphi$	4.7879
$\chi^2_{0.05}$	5.9915

Table 9: *Annualised calibrated Bates parameters.*

Note that the jump parameter, λ, indicates that jumps occur approximately four times per year. The goodness-of-fit statistic suggests that the hypothesis that the observed data are realised from the calibrated Bates SVJ model is not rejected at a 5% level of significance. The Bates SVJ model is, therefore, a plausible data generating model for the FTSE/JSE Top40. Figure 17 compares the density of the calibrated Bates SVJ model to a kernel density estimate of the FTSE/JSE Top40.

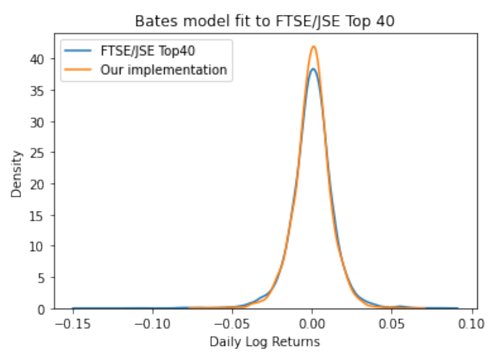

Figure 17: *Comparison of densities under the Bates SVJ model.*

Once again, the two densities align well. Table 10 below compares the first four statistical moments and minimum/maximum values generated by the Bates SVJ model with the daily FTSE/JSE Top40 returns.

Statistic	FTSE/JSE Top40 index	Our implementation
Mean	0.0385%	0.0246%
Std dev	1.3290%	1.1532%
Skewness	-0.4369	-0.2039
Kurtosis	9.4344	4.9364
Minimum	-0.1429	-0.0714
Maximum	0.0845	0.0652

Table 10: *Calibrated Bates model daily statistics for the FTSE/JSE Top40.*

The Bates SVJ model still underestimates the skewness and kurtosis observed in the FTSE/JSE Top40, but improves the fit compared to that of the Heston model.

Given that the Bates SVJ model captures the higher order moments better than the Heston model, we conclude that both stochastic volatility and jumps are required to characterise equity returns in the South African market.

In the next section, we show how the real-world Bates SVJ model can be used in practice by considering a simple volatility targeting strategy.

4 Volatility targeting

In this section, we extend the work of Khuzwayo and Maré [15] by implementing a simulation-based approach to assess the risk and return of various volatility targeting strategies in the South African equity market.

As explained by Khuzwayo and Maré, volatility targeting is an asset allocation strategy that aims to keep the volatility of a portfolio stable by updating the allocation between a risky asset and cash on a regular basis.

Let $\Pi(t)$ denote the time t value of a portfolio consisting of a single risky asset (an equity index) and cash. Mathematically, the value of the portfolio at time $t + dt$ can be written as

$$\Pi(t + dt) = w_S(t)\Pi(t)\left(\frac{S(t + dt)}{S(t)}\right) + w_C(t)\Pi(t)(1 + rdt) + qw_S(t)\Pi(t)(1 + rdt),$$

where r is the rate earned on cash, q is the continuously compounded dividend yield per annum; w_S and w_C are the equity and cash weights given by

$$w_S(t) = \frac{\sigma_{Target}}{\sigma_{Model}(t + 1)}, \quad w_C(t) = 1 - w_S(t).$$

Following Khuzwayo and Maré, we impose the restriction $w_S \leq 1$ so that gearing (borrowing funds to increase equity exposure) is not allowed.

At each time t, we must generate a volatility forecast for time $t + 1$. To do this, we fit a GARCH(1,1) model (see, [7]) to each return series simulated from the Bates SVJ model.

To simplify matters and focus on equity, we use an interest rate of $r = 0\%$ to ignore the effect of interest compounding. For the dividend yield, we set $q = 2.5\%$, roughly representing the long-term dividend yield of the equity market.

Figures 18 to 22 below illustrate a 10% volatility targeting strategy by simulating two paths for the equity index from the Bates SVJ model using an Euler Monte Carlo scheme over a period of one year (252 trading days):

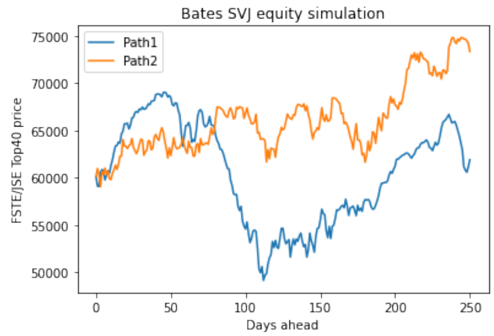

Figure 18: *Bates SVJ equity forecast.*

For each return series simulated from the Bates SVJ model, we fit a GARCH(1,1) model to forecast the volatility. The GARCH(1,1) model is calibrated once every month to a rolling 1000-day history of returns and used to predict 1-week volatility. This is illustrated in Figure 19.

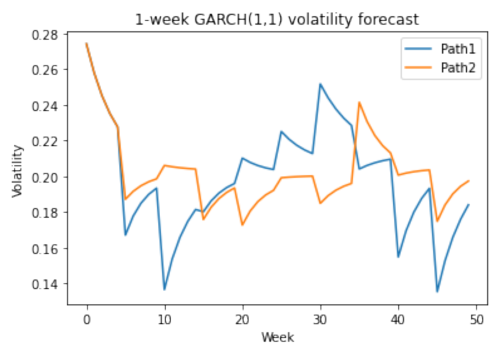

Figure 19: *Weekly volatility forecast.*

At the start of each week, we use the volatility forecast to calculate the equity/cash weights and track the performance of an initial investment of R100.

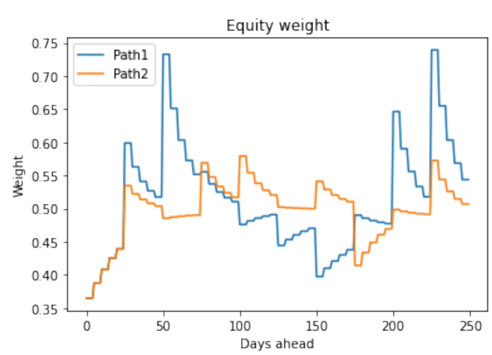

Figure 20: *Equity weight based on 10% volatility target.*

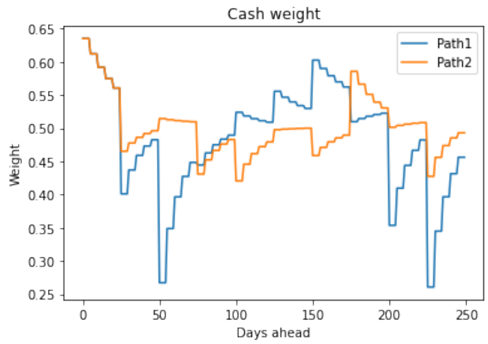

Figure 21: *Cash weight based on 10% volatility target.*

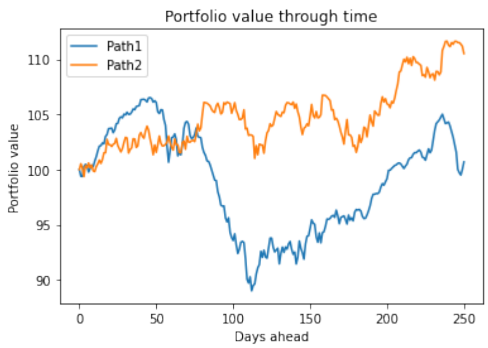

Figure 22: *Portfolio value through time.*

Note that the equity and cash weights are updated at the start of each week and held constant for a one-week period.

Next, we scale the number of simulations to 1000 and test the performance of different volatility targeting strategies over different investment horizons. We consider volatility targets of 10%, 15%, and 20%, and investment horizons of one, three, and five years and compare the performance to that of an equity-only holding strategy. Note that the simulation-based approach allows us to analyse a distribution of returns. The distribution and statistics for each investment horizon and volatility targeting strategy are shown below.

One-year performance

Figures 23 and 24 compare the return and volatility distributions for the various volatility targets and equity-only holding strategy over a 1-year investment horizon.

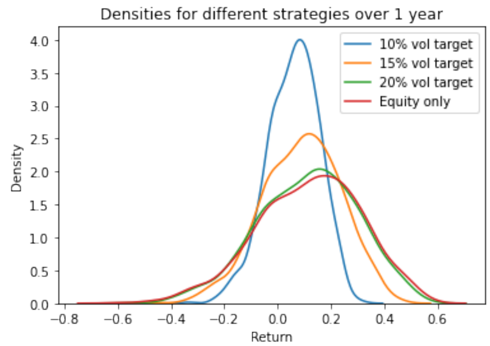

Figure 23: *1-year returns.*

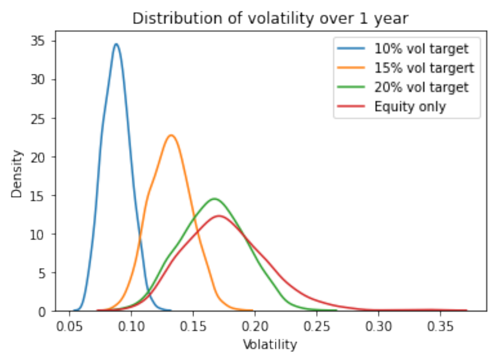

Figure 24: *1-year volatility.*

Table 11 shows the 1-year statistics for the various trading strategies. The results will be discussed at the end of the section.

Statistic	10% vol target	15% vol target	20% vol target	Equity only
Mean return	5.9035%	8.9714%	11.2965%	11.8627%
Mean of volatility	8.8023%	13.1699%	16.6214%	17.6328%
Volatility of volatility	1.0873%	1.6691%	2.6237%	3.3337%
Skewness	-0.3488	-0.2320	-0.2628	-0.3214
Kurtosis	3.0587	2.8589	2.7879	2.9088
Minimum	-33.5001%	-46.0088%	-55.4725%	-60.3300%
Maximum	32.1040%	46.1876%	56.5016%	55.8077%

Table 11: *Statistics for volatility targeting strategies over one year.*

Three-year performance

Figures 25 and 26 compare the return and volatility distributions for the various volatility targets and equity-only holding strategy over a 3-year investment horizon.

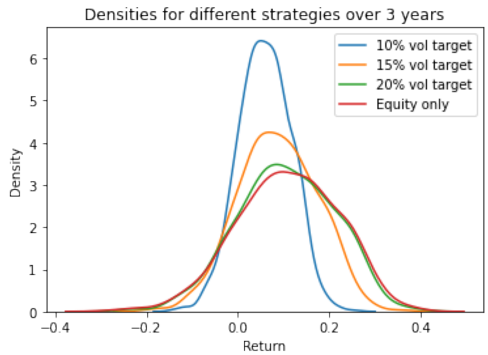

Figure 25: *3-year returns.*

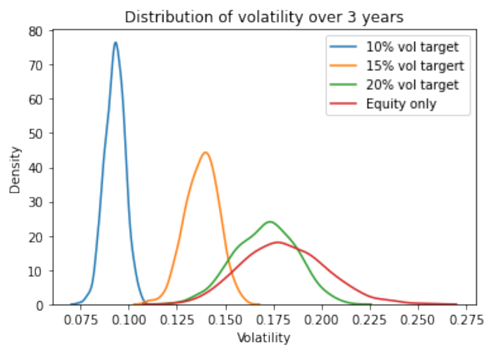

Figure 26: *3-year volatility.*

Table 12 shows the 3-year statistics for the various trading strategies.

Statistic	10% vol target	15% vol target	20% vol target	Equity only
Mean return	5.9788%	8.7532%	10.5564%	10.9141%
Mean of volatility	9.2840%	13.7929%	17.0814%	18.0664%
Volatility of volatility	0.5179%	0.8613%	1.5903%	2.1366%
Skewness	-0.0584	-0.0509	-0.1559	-0.2347
Kurtosis	2.9797	2.8436	2.8178	2.9164
Minimum	-14.1057%	-20.7017%	-25.8108%	-29.3130%
Maximum	25.7383%	36.8075%	41.1602%	41.1486%

Table 12: *Statistics for volatility targeting strategies over three years.*

Five-year performance

Figures 27 and 28 compare the return and volatility distributions for the various volatility targets and equity-only holding strategy over a 5-year investment horizon.

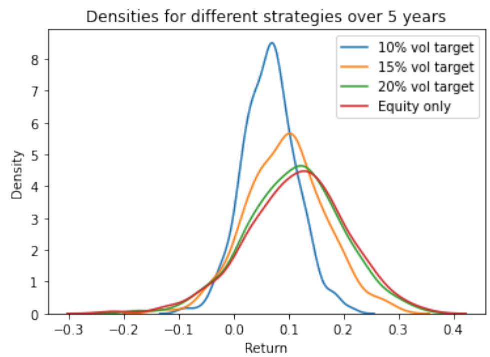

Figure 27: *5-year returns.*

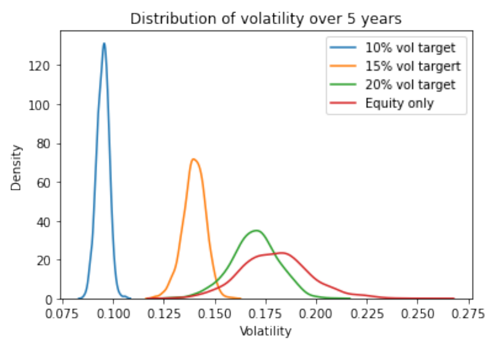

Figure 28: *5-year volatility.*

Table 13 shows the 5-year statistics for the various trading strategies.

Statistic	10% vol target	15% vol target	20% vol target	Equity only
Mean return	6.5284%	9.4335%	11.0569%	11.4293%
Mean of volatility	9.4985%	13.9930%	16.9624%	18.0113%
Volatility of volatility	0.2990%	0.5538%	1.1693%	1.6878%
Skewness	0.0483	0.0028	-0.1638	-0.3107
Kurtosis	3.3085	3.1773	3.1763	3.4514
Minimum	-9.7357%	-14.5351%	-18.8936%	-23.4266%
Maximum	21.8382%	30.0869%	34.7505%	35.3795%

Table 13: *Statistics for volatility targeting strategies over five years.*

Results discussion

There are a couple of interesting observations from the results. Firstly, the mean of the volatility estimate is close to the volatility target for each of the investment horizons. Furthermore, the volatility of volatility estimate decreases as the investment horizon increases. This indicates that the volatility targeting strategy improves for longer investment horizons. Investors can, therefore, expect a targeted volatility with a greater level of certainty for longer investment horizons.

Our results indicate that the volatility target strategy serves to reduce the likelihood of extreme returns and reduces the volatility of volatility.

For all investment horizons, the risk and return of the portfolio increases as the volatility target increases. The 10% volatility target has the lowest risk when viewing the mean of volatility and volatility of volatility estimates, but also the lowest return. On the other hand, an equity only-holding strategy has the highest risk but also the highest expected return. Volatility targeting, therefore, gives investors an effective way of managing the downside risk of a portfolio, but limits the upside potential as shown by the minimum and maximum statistics.

Our findings are consistent with the results in Khuzwayo and Maré [15]. The simulation-based framework introduced in this paper gives investors and fund managers a way of testing portfolio strategies for a wide variety of stressed and benign market conditions.

5 Conclusion

In this paper, we calibrated the Heston and Bates SVJ models to historical S&P500 and FTSE/JSE Top40 returns using the EMM technique of Gallant and Tauchan. First, we confirmed the accuracy of our implementation by comparing our calibrated parameters with Andersen *et al.* [3] by calibrating the Heston and Bates SVJ models over the period 02/01/1953 to 31/12/1996. Our results confirmed that stochastic volatility and jumps are both required to characterise equity returns in the US equity market.

Next, we used the EMM to calibrate the Heston model over different periods to test the stability of the model parameters. Our results suggest that the calibration is sensitive to the input data used and not necessarily stable over time.

The EMM method was then used to calibrate the Heston and Bates SVJ models to FTSE/JSE Top40 returns over the period 30/06/1995 to 31/06/2022. Our calibration results suggest that both stochastic volatility and jumps are required to capture the behaviour of the South African equity market.

The final step was to show a practical application of the Bates SVJ model in the real-world measure. We performed a simulation-based study of various volatility targeting strategies and showed that the risk of a portfolio can be managed effectively by targeting a specific volatility and regularly allocating an investment between a risky asset and cash.

The real-world measure is extremely useful to forecast risk and return. We hope that portfolio managers will find use in our technique to identify future investment opportunities.

References

[1] Aït-Sahalia Y & Jacod J, 2009, *Testing for jumps in a discretely observed process*, Annals of Statistics, 37(1), 184-222.

[2] Aït-Sahalia Y, Jacod J & Li J, 2012, *Testing for jumps in noisy high frequency data*, Journal of Econometrics, 168(2), 207-222.

[3] Andersen T, Benzoni L & Lund J, 2002, *An empirical investigation of continuous-time equity return models*, The Journal of Finance, 57(3), 1239-1284.

[4] Andersen T, Chung H & Sorensen B, 1999, *Efficient method of moments of a stochastic volatility model: A Monte Carlo study*, Journal of Econometrics, 91(1), 61-87.

[5] Bates D, 1996, *Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options*, The Review of Financial Studies, 9(1), 69-107.

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

[7] Brownlees C, Engle R & Kelly B, 2012, *A practical guide to volatility forecasting through calm and storm*, The Journal of Risk, 14(2), 3-22.

[8] Cont, R, 2001, *Empirical properties of asset returns: stylized facts and statistical issues*, Quantitative Finance, 1(2), 223-236.

[9] El Euch O, Gatheral J & Rosenbaum M, 2019, *Roughening Heston*, Risk, 84-89.

[10] Gallant A & Nychka D, 1987, *Semi-nonparametric maximum likelihood estimation*, Econometrica, 55(2), 363-390.

[11] Gallant A & Tauchan G, 1996, *Which moments to match?*, Econometric Theory, 12(4), 657-681.

[12] Gatheral J, 2006, *The volatility surface: A practitioner's guide*, 1st Edition, John Wiley & Sons.

[13] González-Urteaga A, 2012, *Further empirical evidence on stochastic volatility models with jumps in returns*, The Spanish Review of Financial Economics, 10(1), 11-17.

[14] Heston S, 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.

[15] Khuzwayo B & Maré E, 2014, *Aspects of volatility targeting for South African equity investors*, South African Journal of Economic and Management Sciences, 17(5), 691-699.

[16] Olivieri A, Thirurajah S & Ziveyi J, 2022, *Target volatility strategies for group self-annuity portfolios*, ASTIN Bulletin: The Journal of the IAA, 52(2), 591-617.

[17] Poklewski-Koziell W, 2012, *Stochastic volatility models: Calibration, pricing and hedging*, Master's thesis, University of the Witwatersrand, *http://hdl.handle.net/10539/11991*.

[18] Van Dijk M, de Graaf C & Oosterlee C, 2018, *Between $\mathbb{P}$ and $\mathbb{Q}$: The $\mathbb{PQ}$ measure for pricing in asset liability management*, Journal of Risk and Financial Management, 11(4), 67.

[19] Visagie IJH & Grobler GL, 2019, *On the discepancy between the objective and risk neutral densities in the pricing of European options*, ORiON, 35(1), 33-56.