Evaluating Long-Term Performance of Structured Products

Philip Hansen and Mikael Lärfars

Abstract

We investigate the properties of systematic investment vehicles consisting of equity-linked notes in a model with stochastic volatility, random jumps and stochastic interest rate. We consider a setting where the investment horizon is significantly longer than the tenor of the available structured retail products. Long-term asset price trajectories are simulated and performance is evaluated in a quantitative fashion as well as by means of a discretionary scenario analysis.

It is shown that structured products can enhance the risk-return spectrum when introduced in a classical stock-bond mix and that portfolios consisting of multiple structured products reduce the level of timing risk as compared to a so called Single roll. On a more qualitative note, we show that the portfolios of structured products slightly reduce the time-variability of market risk as there is a smoothening on the relative weighting between bonds and options, respectively. Finally, we find that products issued above par, albeit associated with higher risks, show far more attrative return opportunities.

Keywords: Structured products, equity-linked notes, derivatives pricing, portfolio theory

Acknowledgements

We are particularly indebted to our supervisors Karl Lindqvist and William Sjöberg at Nordea Markets, Stockholm and Harald Lang at the Mathematics Department at the Royal Institute of Technology (KTH), Stockholm for insightful comments and guidance. Additionally, the authors kindly thank families and friends for their tireless support.

Stockholm, January 2010

Philip Hansen and Mikael Lärfars

Contents

Chapter 1

Introduction

Generally, structured products can be defined as combinations of elementary financial instruments. In this thesis, we deal with perhaps the most widespread form of structured products within the retail client segment, namely equity-linked notes with principal protection. Such notes typically consist of a zero-coupon bond and a stock option with tenors ranging from two to five years. The special feature in comparison to traditional bonds is that the payoff at maturity is contingent on the performance of the underlying stock or equity index.

Since the first notes denominated in SEK were issued in the late 1980s the market for retail-targeted structured products has gown tremendously. According to Euroclear Sweden, as of July 1, 2009 there were outstanding index- and equity-linked notes demoninated in SEK of approximately SEK 167 billion in nominal terms and the four largest issuers (Nordea, SEB, Svenska Handelsbanken and Swedbank) together issued notes worth of SEK 16.8 billion during the first half of 2009.[1]

This growth serves as evidence of the popularity of structured products and these securities are an easy means of implementing investment strategies based on risk-averse (more specifically, loss-averse) preferences, where the safety of a bond is combined with the opportunities in the equities markets via an option. Additionally, as structured products come in many different forms, with exposure to a variety of asset classes in domestic as well as foreign markets, they generally offer a useful extension and provide increased accessability to the capital markets.

Along with the abovementioned growth, the interest of how a structured product affects the portfolio of a private investor has increased. Several studies (e.g. Goltz, Martellini and Simsek, 2005) have been made investigating how the inclusion of structured products affects an investment portfolio under the assumption that the investment horizon coincides with the tenor of the structured product. From a practical perspective, such assumptions may be too simplistic as an investor's investment horizon may well differ from the tenor of the available structured products and the fact that investors generally have access to a secondary market where long positions in structured products can be liquidated.[2]

In this thesis we investigate the properties of investing into systematic investment vehicles consisting of equity-linked notes or, simply put, portfolios of structured products. More specifically, we construct a class of self-financing, buy-and-hold portfolios involving structured products. We assume the existence of a data-generating process and suggest a general equillibrium model in order to generate plausible future price scenarios and price the according future prices of derivative securities under the risk-neutral measure. The model is calibrated from the Swedish primitive and derivatives markets, respectively.

One of the main reasons for not conducting a historical study is that under the assumption of the existence of a data-generating process we would thus be evaluating a single realization of this process. Further, on back of the relatively short history of the market for structured products we would, at most,

[1] Figures include notes with no principal protection, e.g. so called certificates.

[2] For exceptions, see e.g. Liu and Pan (2003) where no assumptions about the investor's investment horizon are made, although they still do not allow for diversification in tenor.

have a total of two disjunct ten-year periods. These arguments together call for a simulation-based study. Finally, we hope to shed some light on the problem of timing and extreme events and how these concepts are related to portfolio optimization. The problems associated with timing are perhaps best described by the following back-of-the-envelope example: a long position in the OMXS30 index excluding dividends during 1991-2008 roughly increased by 341%, while a similar strategy with the exception of staying out of the stock market (i.e. holding cash) during the worst two years (2002 and 2008) had given a total return of 1135%.

The remainder of this thesis is structured as follows. In Chapter 2 we introduce the model framework followed by a discussion on empirical findings on asset returns and derivatives pricing. Chapter 3 describes the data and the model calibration while Chapter 4 deals with the implementation of our model, more specifically the choice of discretization scheme and simulation of trajectories along with quantitative as well as scenario-based analysis of the portfolios of structured products. We conclude in Chapter 5.

Chapter 2

Modeling financial markets

Most, if not all, of today's continuous-time modeling of financial markets builds on the seminal work of Black and Scholes (1973), here referred to as the Black-Scholes-Merton. One of the major benefits of the Black-Scholes-Merton model compared to more sophisticated expansions is the tractability of the asset price data-generating process as well as the risk-neutral dynamics used for derivatives pricing. Today it is however widely accepted that the assumptions underlying the Black-Scholes-Merton model are too restrictive and it is a well-documented fact that asset returns exhibit both excess kurtosis and skewness which violates its normality assumptions.

Popular attempts to explain these deviations from normality include the introduction of stochastic volatility and random jumps to the asset price process. Several studies investigating asset return characteristics have shown that these models better fit historical data and empirical evidence suggest that models incorporating stochastic volatility and random jumps are indeed well suited to price vanilla options, see e.g. Bakshi, Cao and Chen (1997). A class of model specifications that has received a lot of attention in literature is the class of affine models. Affine models are continuous-time models characterized by drift and variance functions that are linear in risk-factors and include various volatility specifications as well as random jumps in prices and/or volatility. Furthermore, affine models permit closed form solutions on vanilla option prices and in some cases analytical expressions for asset return moments (see e.g. Duffie, Pan and Singleton, 2000). These properties prove helpful as they significantly reduce the complexity of calibration.

2.1 Modeling with a smile in a Hansen-Lärfars framework

In this thesis we simulate long-term trajectories of a stock index, a volatility process and a short rate in order to compute the future value of a portfolio of structured products and compare the long-term performance of this portfolio to other investment alternatives. This approach implies a set of restrictions on our choice of model. Firstly, we require a model that is tractable under the data-generating measure (we will refrain from using the terms objective or real-world measure in order to stress that this is a model framework) as well as under the risk-neutral measure. This is due to the fact that we need to simulate and evaluate returns from an investor's perspective as well as compute option and bond prices in order to find the price processes of the portfolios of structured products. Secondly, we need the data-generating process to capture long-term characteristics of observed market prices whilst the risk-neutral process needs to make accurate out-of-sample predictions on bond and option prices in order for the model to produce realistic price paths. Finally, implementation and calibration to observed data may not be computationally demanding at an unreasonable level. In Sections 2.1.1 and 2.1.2 we define our model of choice and in Section 2.2 we discuss how this model responds to the abovementioned restrictions. We have chosen an affine model with the square-root process of stock index volatility introduced by Heston (1993), random jumps and a Cox, Ingersoll and Ross (1985) model of the short rate.

2.1.1 The data-generating process

We let the stock index return and interest rate to follow data-generating processes on the following form.

$$\begin{aligned} dS_t/S_t &= \left(\mu_0 + r_t - d + \frac{1}{2}V_t\right)dt + \sqrt{V_t}dW_t^S + (J_t - 1)\,dq_t\,(\lambda) - \lambda\mu dt \\ dV_t &= \kappa(\gamma - V_t)dt + \sigma_V\sqrt{V_t}dW_t^V \\ dr_t &= \beta(\alpha - r_t)dt + \sigma_r\sqrt{r_t}dW_t^r \\ dW_t^S dW_t^V &= \rho dt, \end{aligned} \tag{2.1}$$

where S_t, V_t, r_t, and d denote the stock index price (excluding dividend), the variance, the interest rate and the dividend yield, respectively. Further, the jump component $dq_t\,(\lambda)$ is assumed to be a Poisson process with constant intensity λ, independent of the jump size J_t, the driving Wiener-processes dW_t^S, dW_t^V, and dW_t^r while dW_t^S and dW_t^V are assumed to be correlated with correlation ρdt, but uncorrelated with dW_t^r. As Pan (2002), we could model the dividend yield as a stochastic process, but as concluded by Jiang (2002) (on the S&P 500 index) and as seen from the data in Section 3.1.1, relative to the OMXS30 index returns, the daily changes in dividend yield is rather small in magnitude measured by both the mean and standard error. Therefore, for simplicity, we assume a deterministic dividend yield. Furthermore, the jump size is assumed to be log-normally distributed, i.e. $\ln(J_t) \sim N\left(\mu_J, \sigma_J^2\right)$ where $\mu_J = \ln(1+\mu) - \frac{1}{2}\sigma_J^2$. The last term $-\lambda\mu dt$ compensates the change in expected return induced by the jump process. The drift term also consists of μ_0 which describes equity risk premium, including jump risk premium.

Our model of choice is essentially the same model used in Jiang (2002). However, we restrict the volatility risk premium to $V_t/2$ whereas Jiang (2002) only assumes the volatility risk premium to be proportional to V_t. Imposing this restriction simplifies the analytical moments of the equity index log returns and improves the robustness of our calibration procedure. Other affine models with random jumps, stochastic interest rate and stochastic volatility (hereafter the SVJDSI model) include e.g. Pan (2002), Bates (1996, 2000) and Bakshi, Cao and Chen (1997). This is indeed a model with a high degree of complexity and, as Jiang (2002) points out, under certain parameter restrictions it would be equivalent to other popular models. Specifically, with parameter restrictions $\beta = 1$, $\sigma_r = 0$ this is a model with stochastic volatility and random jumps (the SVJD model), with $\beta = \kappa = 1$ $\sigma_r = \sigma_V = 0$ a model with random jumps (the JD model), with $\beta = 1$, $\sigma_r = \lambda = 0$ a model with stochastic volatility (the SV model) and finally with $\beta = \kappa = 1$, $\sigma_r = \sigma_V = \lambda = 0$ this is the Black-Scholes-Merton model (hereafter the BSM model). Popular models that are not special cases of our model include models with random jumps in the variance process, models with random jump intensity, multi-factor volatility models, non-affine models and models that incorporate more sophisticated interest rate dynamics. We will further elaborate on why we refrain from using such extensions in Section 2.2.

2.1.2 The risk-neutral process

As our model of choice is an extension of the BSM model with additional risk factors but without the introduction of additional risky assets we are modeling a non-complete market and, hence, the risk-neutral measure is not unique. We will below state the necessary assumptions for a risk-neutral process on a similar form as Equation 2.1.

Given the data-generating process stated in Equation 2.1 with deterministic jump intensity and under the assumption that volatility risk premium for the stock index and interest rate processes, Φ_V and Φ_r' are stochastic and proportional to the corresponding variance processes, i.e. on the form $\Phi_V = \xi V_t$, $\Phi_r = \zeta r_t$, Jiang (2002) uses the results of Bates (1988) and Cox, Ingersoll and Ross (1985) to show that the corresponding risk-neutral process is described by Equation 2.2.[1]

[1]See Jiang (2002) for an explicit expression of the pricing kernel.

$$
\begin{aligned}
dS_t/S_t &= (r_t - d)\,dt + \sqrt{V_t} dW_t^{S*} + (J_t^* - 1)\,dq_t^*\,(\lambda^*) - \lambda^*\mu^* dt \\
dV_t &= (\kappa\,(\gamma - V_t) + \Phi_V)\,dt + \sigma_V \sqrt{V_t} dW_t^{V*} \\
dr_t &= (\beta\,(\alpha - r_t) + \Phi_r)\,dt + \sigma_r \sqrt{r_t} dW_t^{r*} \\
dW_t^{S*} dW_t^{V*} &= \rho dt.
\end{aligned}
\tag{2.2}
$$

Bates (1996) has shown that the above linear assumptions on risk premia is equivalent to assuming a market log utility function of terminal wealth. We will however leave the discussion on utility theory and implicitly motivate the choice of pricing kernel by discussing how the data-generating process fits observed stock index returns and the pricing performance of the risk-neutral process.

2.2 Model properties – a tale of heavy tails

In this section we cover some important empirical and theoretical findings in previous studies that elaborate on asset return properties, derivatives pricing and different model specification. In particular, we use previous findings in the abovementioned areas to motivate our choice of model. Furthermore, we utilize previous work on affine models to derive some pricing formulas as well as some statistical properties of the SVJDSI model.

2.2.1 Equity index return properties

There are several aspects of asset return characteristics that are not captured by simpler models. Cont (2001) presents an extensive overview of well-documented asset return properties, e.g. volatility clustering, negative skewness, excess kurtosis and negative correlation with volatility. The negative correlation with volatility is usually described as a leverage effect, that is a decrease in stock prices increases debt/equity ratios and, hence, leverage and volatility is thereby increased. In our model, this leverage effect would be captured by a negative value of the correlation coefficient ρ.

These properties are also studied by Chernov et al. (2003) who calibrate several model specifications, including the most common affine models, to observed data and evaluate their ability to explain asset return properties. In particular, it is concluded that a one-factor stochastic volatility model without jumps is insufficient if one wants to capture both volatility clustering and tail behavior and, hence, the model is rejected. Moreover, the authors conclude that modeling tail behavior by including random jumps allow for the one-factor volatility to describe volatility persistence more accurately. This is due to the fact that accurate modeling of tail behavior with only stochastic volatility requires a higher speed of mean reversion which reduces volatility clustering. We have found little evidence of two-factor volatility, with or without jumps, being superior to a one-factor volatility, random jump model. This, combined with the additional amount of parameters and loss in degrees of freedom induced by a two-factor model, implies that our model specification should be sufficient to describe the abovementioned asset properties. Andersen, Benzoni and Lund (2002) reach similar results as do Chernov et al. (2003) and the auhors reject models that exclude random jumps and/or stochastic volatility. Furthermore, they show that state dependent jump intensity does not provide any additional explanatory power over constant jump intensity. They do, however, discuss the possibility that this result could stem from estimation issues. Regardless, we have not found sufficient proof of additional explanatory power in stochastic jump density to motivate the additional computational complexity it would incur.

2.2.2 Pricing power in the interests of rates and options

It is well documented that the BSM model consistently misprice options rendering an implied volatility smile or skew and a non-flat term structure of volatility. In particular, implied volatilities are generally decreasing in strike price inferring that the BSM model underestimates the risk of large negative price shocks. The studies of models that better fit observed market prices and price out-of-sample options more accurately are numerous. Generally, it appears the convention is to re-calibrate models on a daily basis.[2] This is of course internally inconsistent as parameters assumed to be constant in model specifications are continuously changing. As our study uses simulated price paths we can not use this method of calibration but must trust the model calibration to be robust over time. The major drawback is that we consequently require a high level of out-of-sample pricing power.

Similar to the conclusions drawn on the data-generating process in the Section 2.2.1 Bakshi, Cao and Chen (1997) conclude that stochastic volatility or random jumps alone does not capture the risks priced by index options markets. In particular, stochastic volatility fails to capture short-term kurtosis and hence misprice short term options while pure jump diffusion models do not fit the prices of options with longer tenors. For a good out-of-sample cross-sectional fit they recommend including both stochastic volatility and random jumps.[3] In contrast, stochastic interest rates does not reduce cross-sectional biases but the authors recommend implementing the same in order to achieve a higher overall pricing precision. Jiang (2002) does, however, argue that the improvement from stochastic interest rates is negligible. Regardless, our main motivation for including stochastic interest rates is the need to simulate reasonably realistic investment scenarios and the possibility to price bonds from a non-flat, time-varying yield curve.

Pan (2002) shows that including state-dependent jump intensity significantly improves the fit to observed option prices. This result is however based on a comparison with models excluding jumps, i.e. these results do not imply that stochastic jump density describes the options markets more accurately than constant jump density. Further, Eraker (2004) concludes that including jumps in the variance process has a significant positive impact on out-of-sample pricing performance. Although this is indeed a reasonable extension we choose not to include jumps in volatility as it would incur additional complexity in the calculation of asset return moments used for GMM estimation. In particluar, we need to know the explicit distribution of the variance process in order to compute the unconditional characteristic function of the stock index log returns.

The major weakness of our proposed model appears to be the interest rate dynamics. It has indeed been shown that the Cox-Ingersoll-Ross model fails to capture some well-observed interest rate charachteristics. Among others, Andersen, Benzoni and Lund (2004) conclude that stochastic volatility is crucial for the interest rate model to fit observed U.S. short rates and that random jumps helps explaining outliers in their data set.

As a concluding remark, it should be noted that studies where both stock price and interest rate dynamics are modelled generally aim at pricing equity derivatives. Since the improvement in pricing ability from more advanced interest rate models is negligible as compared to the improvements from more sophisticated stock price dynamics this has resulted in a lack of research on models encompassing both stock price dynamics and more sophisticated interest rate modeling. Consequently, we have not come across any research on how to price equity options in such a setting. Additionally, as with the variance process, we need to know the explicit distribution of the interest rate in order to compute the unconditional properties of the stock index log returns. Hence, we choose not to extend the model with more advanced interest rate dynamics.

2.2.3 Conditional and unconditional properties

To be able to use the calibration method described in Section 3.1 we need to derive the unconditional moments of the stock index log returns and the conditional moments of the interest rate. We follow the procedure of Jiang (2002) and solve the Kolmogorov backward equation for the conditional joint characteristic function of $\Delta \ln S_t$, r_t and V_t, where $\Delta \ln S_t = \ln S_t - \ln S_{t-\Delta}$. Since the charachteristic functions, and, hence, the analytical moments, presented by Jiang (2002) contain some errors, we have chosen to include a full derivation in Appendix A.

[2] See e.g. Bates (1996) or Bakshi, Cao and Chen (1997)

[3] The term cross-sectional refers to a model being able to accurately price options varying across tenors as well as strike prices.

$$
\begin{aligned}
&\psi\left(\Delta \ln S_{t+\Delta}, V_{t+\Delta}, r_{t+\Delta}; \phi_1, \phi_2, \phi_3 | V_t, r_t\right) = \\
&\mathrm{E}\left[\exp\left(i\phi_1 \Delta \ln S_{t+\Delta} + i\phi_2 V_{t+\Delta} + i\phi_3 r_{t+\Delta}\right) | V_t, r_t\right] = \\
&\exp\left(C\left(\phi_1, \phi_2, \phi_3, \Delta\right) + \left(i\phi_2 + D\left(\phi_1, \phi_2, \Delta\right)\right) V_t + \left(i\phi_3 + B\left(\phi_1, \phi_3, \Delta\right)\right) r_t\right) \\
&\exp\left(\Delta\lambda\left(e^{i\phi_1\mu_J - 1/2\phi_1^2\sigma_J^2} - 1\right)\right),
\end{aligned} \tag{2.3}
$$

where $C\left(\phi_1, \phi_2, \phi_3, \Delta\right)$, $D\left(\phi_1, \phi_2, \Delta\right)$ and $B\left(\phi_1, \phi_3, \Delta\right)$ are given in Appendix A.

Using Equation 2.3, that V_t follows a Gamma distribution with density function $f_{V_t}\left(x\right) = \frac{\theta^p}{\Gamma(p)} x^{p-1} e^{-\theta x}$, where $\theta = \frac{2\kappa}{\sigma_v^2}$, $p = \frac{2\kappa\gamma}{\sigma_v^2}$ and that r_t follows a Gamma distribution with density function $f_{r_t}\left(x\right) = \frac{\theta^p}{\Gamma(p)} x^{p-1} e^{-\theta x}$, where $\theta = \frac{2\beta}{\sigma_r^2}$, $p = \frac{2\beta\alpha}{\sigma_v^2}$ we can compute the unconditional charachteristic functions defined in Equations 2.4 and 2.5.

$$
\begin{aligned}
&\psi\left(\Delta \ln S_{t+\Delta}, r_{t+\Delta}; \phi_1, \phi_3\right) = \\
&\mathrm{E}\left[\exp\left(i\phi_1 \Delta \ln S_{t+\Delta} + i\phi_3 r_{t+\Delta}\right)\right] = \\
&\exp\left(C\left(\phi_1, 0, \phi_3, \Delta\right) - \frac{2\kappa\gamma}{\sigma_V^2} \ln\left(1 - \frac{\sigma_V^2 D\left(\phi_1, 0, \Delta\right)}{2\kappa}\right) - \frac{2\beta\alpha}{\sigma_r^2} \ln\left(1 - \frac{\sigma_r^2 B\left(\phi_1, \phi_3, \Delta\right)}{2\beta}\right)\right) \\
&\exp\left(\Delta\lambda\left(e^{i\phi_1\mu_J - 1/2\phi_1^2\sigma_J^2} - 1\right)\right)
\end{aligned} \tag{2.4}
$$

$$
\begin{aligned}
&\psi\left(\Delta \ln S_{t+\tau+\Delta}, \ln S_{t+\Delta}; \varphi_1, \varphi_2\right) = \mathrm{E}\left[\exp\left(i\varphi_1 \Delta S_{t+\tau+\Delta} + i\varphi_2 S_{t+\Delta}\right)\right] = \\
&\exp\left(C\left(\varphi_1, 0, 0, \Delta\right) + C\left(0, -iD\left(\varphi_1, 0, \Delta\right), -iB\left(\varphi_1, 0, \Delta\right), t - \Delta\right) + C\left(\varphi_2, -iD^*, -iB^*, \Delta\right)\right) \\
&\left(1 - \frac{\sigma_V^2\left(D^* + D\left(\varphi_2, -iD^*, \Delta\right)\right)}{2\kappa}\right)^{2\kappa\gamma/\sigma_V^2} \left(1 - \frac{\sigma_r^2\left(B^* + B\left(\varphi_2, -iB^*, \Delta\right)\right)}{2\beta}\right)^{2\beta\alpha/\sigma_r^2} \\
&\exp\left(\Delta\lambda\left(e^{i\varphi_1\mu_J - 1/2\varphi_1^2\sigma_J^2} - 1\right) + \Delta\lambda\left(e^{i\varphi_2\mu_J - 1/2\varphi_2^2\sigma_J^2} - 1\right)\right)
\end{aligned} \tag{2.5}
$$

Now, by differentiating Equations 2.3, 2.4 and 2.5 we can compute the necessary moments described in Section 3.1.3. Refer to Appendix B for a Monte Carlo verification of the exact moments.

2.2.4 Risk-neutral properties

As Cox, Ingersoll and Ross (1985) show, given the risk-neutral representation in Equation 2.2, the price of a zero-coupon bond $B\left(t, \tau\right)$ at time t with tenor τ is on the form as described in Equation 2.6.

$$
B\left(t, \tau\right) = a\left(t, \tau\right) e^{-b(t,\tau) r_t}, \tag{2.6}
$$

where $a\left(t, \tau\right) = \left(\frac{2\Gamma e^{(\beta^* + \Gamma)\tau \backslash 2}}{(\beta^* + \Gamma)(e^{\Gamma\tau} - 1) + 2\Gamma}\right)^{2\gamma_r/\sigma_r^2}$, $b\left(t, \tau\right) = \frac{2e^{\Gamma\tau} - 1}{(\beta^* + \Gamma)(e^{\Gamma\tau} - 1) + 2\Gamma}$, $\Gamma = \sqrt{\beta^* + 2\sigma_r^2}$, $\gamma_r = \beta\alpha$, $\beta^* = \beta - \zeta$. Furthermore, Jiang (2002) uses the Fourier inversion technique propoesed by Heston (1993) to show that 2.2 implies european call option prices $C\left(t, \tau, S_t, K, r_t, V_t\right)$ on the following form. Similar formulais are also derived by e.g. Bates (1996), Scott (1997) and Bakshi, Cao, Chen (1997).

$$
C\left(t, \tau, S_t, K, r_t, V_t\right) = S_t \Pi_1\left(t, \tau, S_t, K, r_t, V_t\right) - K B\left(t, \tau\right) \Pi_2\left(t, \tau, S_t, K, r_t, V_t\right), \tag{2.7}
$$

where $\Pi_j\left(t, \tau, S_t, K, r_t, V_t\right) = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left(\frac{e^{-i\phi \ln K} f_j(t,\tau,S_t,K,r_t,V_t,\phi)}{i\phi}\right) d\phi$ while f_1 and f_2 are given in Appendix C.

Chapter 3

Model calibration

We estimate the data-generating and risk-neutral processes, specified in Equation 2.1 and 2.2, respectively, using information from both observations of the underlying asset returns and the options market. Following Jiang (2002) we employ a two-step estimation procedure where, in the first step, the underlying data-generating model is estimated from asset return observations. More specifically, we start out by estimating the interest rate process parameters and these estimates are then utilized when estimating the stock return process parameters, as has been done by e.g. Pan (2002). These parameters are estimated using Generalized Method of Moments (GMM), described in detail below. In the second step the preference-related parameters are estimated from options market data. As mentioned by Jiang (2002), this implementation is relatively easy and it suits our needs as we employ the data-generation and risk-neutral processes for simulation and derivatives pricing purposes, respectively.

3.1 The data-generating process

Estimation of non-linear latent variable models is by no means a trivial task as the stochastic volatility is unobservable and, hence, the model can not be estimated using standard Maximum Likelihood Estimation (MLE). Various estimation methods for stochastic volatility models have been proposed over the past two decades, yet most of these are simulation-based and very computationally intensive (Jiang, 2002). We exploit that the characteristic functions of the asset returns and joint asset returns can be derived analytically and, thus, the exact moments of asset returns are available. That is, the calibration procedure we employ is based on exact moments of the continuous-time process rather than stemming from a discrete approximation.

3.1.1 Swedish market data

For the estimation of the data-generating process parameters we use daily Swedish market data of the large cap price index OMXS30 from March 28, 1991 until March 31, 2009. As to the theoretical concept of the short rate, we use the 3-month Swedish treasury bill (Sw. statsskuldväxel). This is a necessary compromise between literally using the interest rate with the shortest tenor available (e.g. overnight rates) and avoiding some of the associated microstructure effects (Jiang, 2002). Again, we use daily market data from March 28, 1991 until March 31, 2009.

The continuous dividend yield is estimated using daily observation of the OMXS30 price and gross index, respectively. In an attempt to reduce the dependence structure, the continuous dividend yield, d, is estimated via the daily differences between the log returns of the gross and price indexes using data from January 2, 2002 until October 31, 2008. Figure 3.1 shows charts of the daily quotes of OMXS30 price index and the SSVX 3M interest rates, whereas Figure 3.2 shows a comparison of the OMXS30 price and gross indexes, respectively, along with our estimate of the continuous dividend yield. Finally, Table 3.1.1 shows summary data statistics of the historical Swedish market data used in the calibration procedure described below.

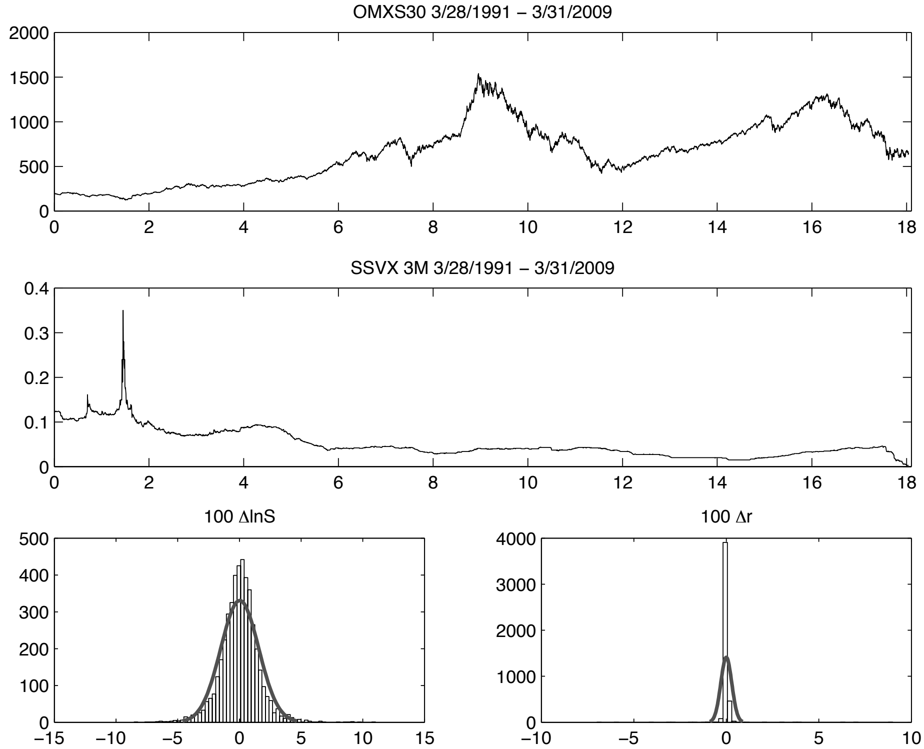

Figure 3.1: Upper: Historical data of the OMXS30 price index from March 28, 1991 until March 31, 2009. Middle: Historical data of the SSVX 3M yield during the same time period. Lower: Historical daily changes as specified of the OMXS30 price index and the SSVX 3M yield.

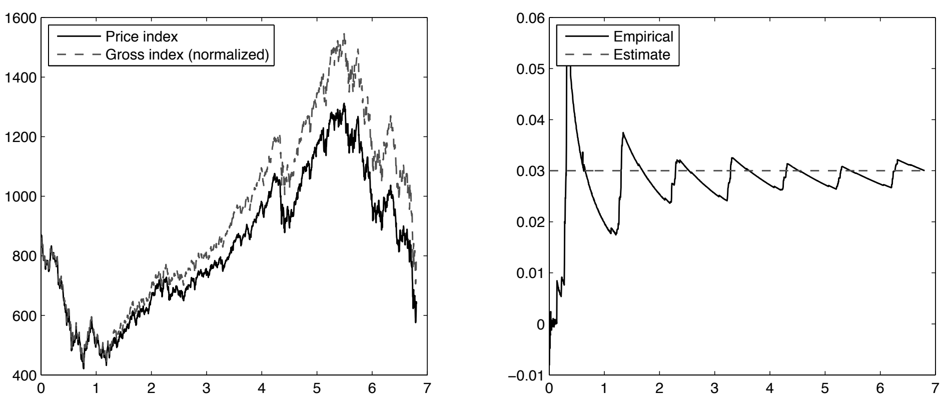

Figure 3.2: Left: Historical data of the OMXS30 price and gross indexes from January 2, 2002 until October 31, 2008. Right: Historical and estimated continuous dividend yield, i.e. $d = t^{-1}\ln(G_t/S_t)$, for the same time period.

Static properties

	N	Mean	St. dev.	Skewness	Kurtosis	Min	Max
$100 \times \Delta \ln S_{t_i}$	4510	0.027	1.541	0.164	6.924	-8.527	11.023
$100 \times \Delta r_{t_i}$	4510	-0.003	0.295	6.210	455.174	-7.000	9.000

Dynamic properties

	N	$\rho(1)$	$\rho(2)$	$\rho(3)$	$\rho(4)$	$\rho(5)$	$\rho(10)$	$\rho(20)$
$100 \times \Delta \ln S_{t_i}$	4510	0.017	-0.041	-0.036	0.013	-0.011	-0.022	-0.000
$(100 \times \Delta \ln S_{t_i})^2$	4510	0.185	0.206	0.205	0.154	0.210	0.180	0.110
$100 \times \Delta r_{t_i}$	4510	0.005	-0.339	-0.115	0.210	0.227	-0.120	-0.009

Table 3.1: Summary data statistics of the OMXS30 price index daily log returns and the SSVX 3M yield changes from March 28, 1991 until March 31, 2009. Autocorrelation with lag l is denoted $\rho(l)$ while $\Delta \ln S_{t_i} = \ln S_{t_i} - \ln S_{t_{i-1}}$ and $\Delta r_{t_i} = r_{t_i} - r_{t_{i-1}}$ with $i = 2, \ldots, N$.

3.1.2 Generalized Method of Moments

Generalized Method of Moments (GMM) is an econometric procedure for estimating the parameters of a given model. Hansen (1982) developed GMM as an extension to the Classical Method of Moments estimators, the latter dating back more than a century. The basic idea of GMM is to choose parameters so as to match the moments of the model to those of the data as closely as possible. The moment conditions are chosen under the implementer's discretion based on the problem at hand, which itself serves as a proof of the generality of the method. A weighting matrix determines the relative importance of matching each moment. Most common estimation procedures can be couched in this framework, including OLS, 2SLS and in some cases even MLE (Cliff, 2003).

The first key advantage of GMM over other estimation procedures is that the initial/underlying statistical assumptions of stationarity and ergodicity are relatively weak when comparing to the more traditional assumption that the data are independent and identically-distributed. Further, unlike MLE, GMM does not put distributional assumptions on the data, although the GMM moment conditions are indeed functionally parametric. Among the more obvious drawbacks is a loss of efficiency over methods such as MLE. Thus, GMM offers a compromise between the efficiency of MLE and robustness to deviations from e.g. normality (Arnold and Crack, 1999).

3.1.3 Implementation and estimation results

The first stage of GMM is to construct the so called population moments $f_t(\theta)$ such that the expectation of the moment vector is equal to zero, i.e. $\mathrm{E}[f_t(\theta)] = 0$. For the interest rate process, we let

$$\varepsilon_t^r = r_t - \mathrm{E}[r_t | r_{t-\Delta}],$$

$$t = 2, 3, \ldots, T.$$

be the de-meaned interest rate process. The expectations of ε_t^r are calculated exactly as in Section 2.2.3. We use the the first two conditional moments with the lagged variable as instrumental variable. These are the same moment conditions used by Chan, Karolyi, Longstaff and Sanders (1992), only that ours are exact as they are derived from the continuous-time model. The moment conditions are formally stated in Equation 3.1.

$$f_t^r(\theta) = \begin{bmatrix} \varepsilon_t^r - \mathrm{E}[\varepsilon_t^r] \\ (\varepsilon_t^r)^2 - \mathrm{E}[(\varepsilon_t^r)^2] \end{bmatrix} \otimes \begin{bmatrix} 1 \\ r_{t-\Delta} \end{bmatrix} = \begin{bmatrix} \varepsilon_t^r - \mathrm{E}[\varepsilon_t^r] \\ (\varepsilon_t^r - \mathrm{E}[\varepsilon_t^r])\, r_{t-\Delta} \\ (\varepsilon_t^r)^2 - \mathrm{E}\Big[(\varepsilon_t^r)^2\Big] \\ \Big((\varepsilon_t^r)^2 - \mathrm{E}\Big[(\varepsilon_t^r)^2\Big]\Big)\, r_{t-\Delta} \end{bmatrix} \tag{3.1}$$

$$t = 2, 3, \ldots, T.$$

Similarly to the interest rate process, we let

$$\varepsilon_t^S = \Delta\text{ln}S_t - \text{E}[\Delta\text{ln}S_t],$$

be the de-meaned asset return process, where $\Delta \ln S_t = \ln S_t - \ln S_{t-\Delta}$. Again, the expectations of ε_t^S are calculated exactly as in Section 2.2.3. There are obviously infinitely many moments that may be included in the GMM estimation. When determining the number of moments used in the estimation a fundamental trade-off applies, namely that the inclusion of additional moments improves estimation performance for a given degree of precision in the estimation of the weighting matrix, but in finite samples this must be balanced against the deterioration in the estimate of the weighting matrix as the number of moments increases (Jiang, 2002). Very high order moments should be avoided due to their erratic finite sample behavior caused by the presence of fat tails in the asset return distribution. Hence, we move our attention to the lower order moments, which is consistent with Jiang (2002), Andersen and Sørensen (1996) and Jacquier, Polson and Rossi (1994). The asset return moment conditions are described in Equation 3.2.

$$f_t^S(\theta) = \begin{bmatrix} \left(\varepsilon_t^S\right)^k - \text{E}\left[\left(\varepsilon_t^S\right)^k\right] \\ \left(\varepsilon_t^S\right)^2\left(\varepsilon_{t-\tau}^S\right)^2 - \text{E}\left[\left(\varepsilon_t^S\right)^2\left(\varepsilon_{t-\tau}^S\right)^2\right] \end{bmatrix} \tag{3.2}$$

$$k = 1, 2, 3, 4, 5;$$

$$\tau = 1, 2, 3, 4, 5;$$

$$t = 7, 8, \ldots, T.$$

Further, the jump component affects only the unconditional moments, thus the first group of moment conditions in Equation 3.2 is important for the estimation of jump parameters. Since stochastic volatility and random jump both allow for skewness and excess kurtosis, it is imperative to include the fifth moment for the estimation of jump parameters and, accordingly, we set $k = 1, 2, 3, 4, 5$. Secondly, the autocorrelation of squared asset return is determined by the dynamics of the stochastic volatility process and its correlation with asset returns. Thus, the second group of moment conditions in Equation 3.2 is important for the identification of the volatility risk premium and volatility dynamics. As to the second group of moment conditions, we follow Jiang (2002) who proposes using different lags, namely $\tau = 1, 2, 3, 4, 5$ on back of empirical evidence suggesting that autocorrelation is varying over time.

The GMM sample moments, that is the natural sample counterpart of the population moments, are then defined as

$$g_T(\theta) = \frac{1}{T}\sum_{t=1}^{T} f_t(\theta), \tag{3.3}$$

where T is the number of available observations or sample size.[1] Further, the GMM objective function is defined as

$$J_T = g_T' W_T g_T, \tag{3.4}$$

where W_T is a positive-definite weighting matrix. We can see that GMM is a minimum distance estimator, i.e. we set the weighted sum of squared sample moments as close to zero as possible. The weighting matrix instructs on how much attention to pay to each moment. Hence, parameter estimates $\hat{\theta}$ are found by solving

$$\hat{\theta} = \arg \min_{\theta \in \Theta} J_T. \tag{3.5}$$

Additionally, GMM offers an overall test of the model by testing whether the "extra" sample moments are sufficiently close to zero relative to their distribution. Under the null hypothesis that the model is true, the minimized value of $J_T(\theta)$ in 3.4 is χ^2-distributed with degrees of freedom equal to the number of orthogonality conditions net of the number of parameters to be estimated. This χ^2-statistic provides a

[1] As a result of Hansen's (1982) seminal work, notations as in e.g. Equation 3.3 where T denotes sample size have become convention, although a perhaps more appropriate notation here would include t_i, where $i = 1, \ldots, N$ and N is the sample size.

goodness-of-fit test for the model, where a high statistic suggests that the model is misspecified. Formally put, the test statistic is asymptotically χ^2-distributed as follows

$$T \times J_T(\hat{\theta}) = T \times g_T'(\hat{\theta}) W_T g_T(\hat{\theta}) \xrightarrow{d} \chi^2_{m-p}, \tag{3.6}$$

where m is the number of moment conditions and p is the number of parameters.

Moments are chosen in such a fashion that the estimation is over-identified, i.e. there are more moment conditions than there are parameters to be estimated. To obtain asymptotically efficient estimates we set $W_T = \widehat{S}^{-1}$, where $\widehat{S}$ is an estimate of the spectral density matrix of population moment functions. This choice of the weighting matrix (sometimes referred to as efficient GMM) secures the smallest asymptotic covariance matrix of the vector of estimated parameters, $\hat{\theta}$ (Hansen, 1982). This spectral density or long-run covariance matrix is defined

$$S = \sum_{j=-\infty}^{\infty} \mathrm{E}[f_t(\theta) f_{t-j}(\theta)']. \tag{3.7}$$

The spectral density matrix allows for serial correlation and heteroskedasticity in the observations of the moments function. We utilize the perhaps most popular, consistent estimate of the spectral density matrix, namely the one proposed by Newey and West (1987). The Newey-West estimator is defined as

$$\widehat{S} = \widehat{S}_0 + \sum_{j=1}^{k} \left(1 - \frac{j}{k+1}\right) \left(\widehat{S}_j + \widehat{S}_j'\right), \tag{3.8}$$

where

$$\widehat{S}_j = \frac{1}{T} \sum_{t=j+1}^{T} f_t(\theta) f_{t-j}(\theta)',$$

and k is the number of lags. Note that when the number of lags are set to 0 the spectral density matrix collapses to an ordinary sample covariance matrix, a procedure relied upon under the assumption that each row in the population moments vector can be considered to consist of i.i.d. variables.

In the case of an over-identified estimation, GMM is a two-stage estimator. First, we minimize Equation 3.4 using the identity matrix as weighting matrix, i.e. $W_T = I$. This means that we consider all moments equally important. We then insert the estimated parameter vector $\hat{\theta}$ into Equation 3.8 and inverse to get W_T. Second, we minimize Equation 3.4 again, only this time using the weighting matrix, W_T, from the previous step. We re-estimate the model parameters using increasing lag lengths until the lag length has negligible effect on the prevailing value of the objective function, J_T. Arnold and Crack (1999) propose $m \approx \sqrt{T}$ as a rule of thumb, whereas several other (e.g. Cliff, 2003) propose a smaller lag length ($m \approx \sqrt[3]{T}$). We use a lag length of $m = 20$ in our estimation.

The resulting parameters for the interest rate and asset return processes are presented in Table 3.2 and Table 3.3, respectively. Firstly, we conclude that the BSMSI provides very little explanatory power, whereas the introduction of stochastic volatility proves to be a great improvement. Secondly, while it seems that the jump diffusion component provides no additional explanatory power to the data-generating process we justify its existence on back of the fact that removing the corresponding jump risk premium severly deteriorates the risk-neutral pricing power. However, when introducing both stochastic volatility and jump diffusion we see that both the speed of mean reversion in volatility as well as the jump intensity are reduced, hence supporting a volatility clustering effect.

	SI
α	0.0316
β	0.430
σ_r	0.0688
$T \times J_T$	1.4
p-value	24.12%
d.f.	1

Table 3.2: Prevailing parameter estimates of the interest rate process.

	BSMSI	SVSI	JDSI	SVJDSI
μ_0	0.0968	0.0715	0.0798	0.0704
$\sqrt{\gamma}$	0.207	0.232	0.190	0.226
κ	-	2.85	-	1.02
σ_V	-	0.517	-	0.340
ρ	-	-0.118	-	-0.235
μ_J	-	-	-0.00381	-0.00763
λ	-	-	5.47	2.71
σ_J	-	-	0.0377	0.0405
$T \times J_T$	24.6	4.5	22.2	3.0
p-value	0.19%	47.49%	0.05%	22.38%
d.f.	8	5	5	2

Table 3.3: Prevailing parameter estimates of the stock index process.

3.2 The risk-neutral process

In order to estimate the preference-related parameters β^*, κ^*, μ^*, and λ^* we employ an OLS-type calibration procedure, where model-implied options and bond prices are matched to observed market prices. As mentioned in Section 2.1 we do not utilize daily recalibration and hence need our parameter estimates to be robust over time. Considering this we use a seven month data window, i.e. the period from April 30, 2009 until October 30, 2009. Note that this period does not overlap the dataset used in the calibration of the data-generating process. Since the calibration procedure involves optimization over prices that must be retrieved by means of numerical integration the process is indeed time-consuming and we settle for a low data frequency. The dataset consists of two liquid options and three government bonds per month, all quoted at the last trading day of the month.

The estimation results could likely be improved further if we could control the time synchronization between stock index value and option prices. To adress this problem we use liquid options and the average of the bid and offer quotes, rather than the closing price. Moreover, Jiang (2002) concludes that for options on the S&P 500 index, parameter estimates tend to vary over moneyness. Hence, we use liquid "near-the-money" options since the price of these options will have the greatest impact on the performance of the structured products in our study.[2]

We use a minimum squared error loss function for both the bond and the option calibration. This loss function appears to be the most commonly used, although it allows for expensive options to have a greater impact on the calibration results. Westermark (2009) does, however, conclude that the choice of loss function has limited impact on out-of-sample pricing and, hence, other loss functions would are likely to work as well. The results are found in Table 3.4.

Bond calibration	
β^*	0.216
MSE	SEK 0.32
MPE	0.48%
Option calibration	
κ^*	11.1
μ^*	-0.0435
λ^*	4.83
MSE	SEK 2.3
MPE	1.5 %

Table 3.4: Prevailing estimates of the risk-neutral parameters. MSE denotes mean squared pricing error and MPE denotes absolute mean percentage pricing error

As discussed in Section 2.2 and as can be seen in Figure 3.3 the model bond prices lack some distinct features observed in the market. In particular, the model fails to capture the convex yield curve observed e.g. in October 2009. This problem should stem from the model specification rather than from the parameter estimates and we still find the model an improvement over constant interest rates which would generate flat, state-independent, yield curves.

[2] In particular, see Section 4.2.1 for a discussion on e.g. participation rate and how it relates to moneyness.

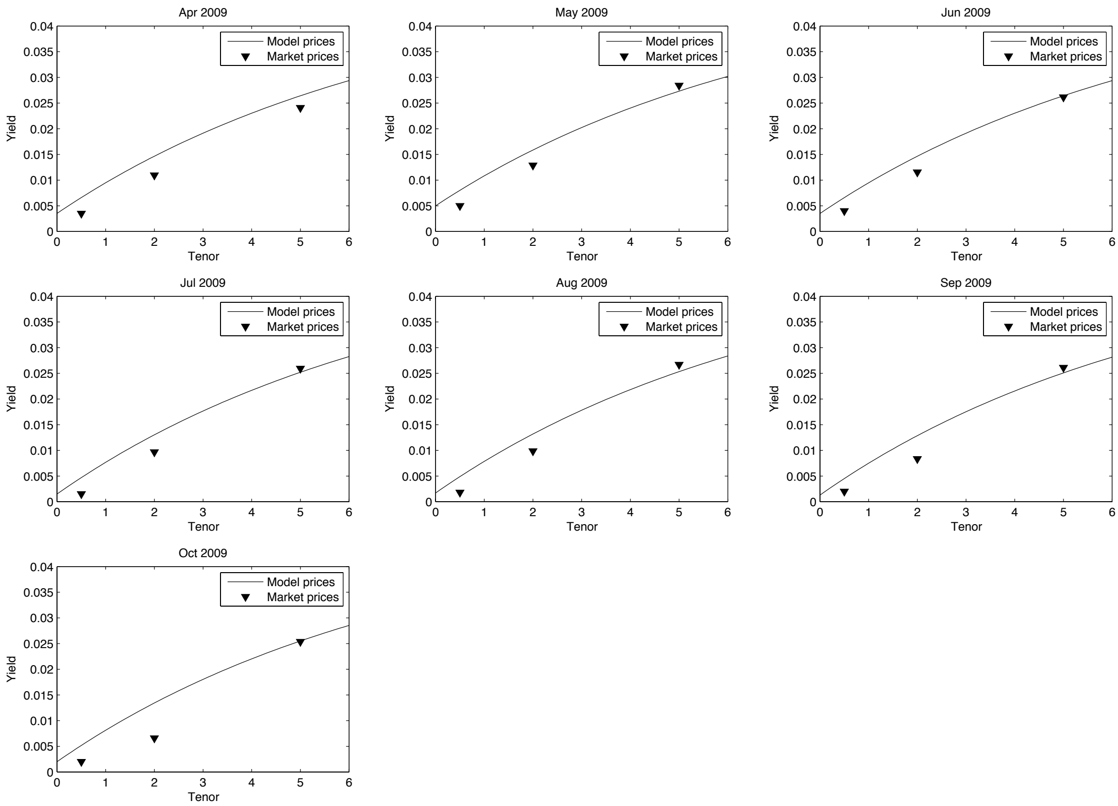

Figure 3.3: Plots of model and market prices (or yields) of the Swedish 6-month treasury bill and the Swedish 2- and 5-year government bonds, respectively.

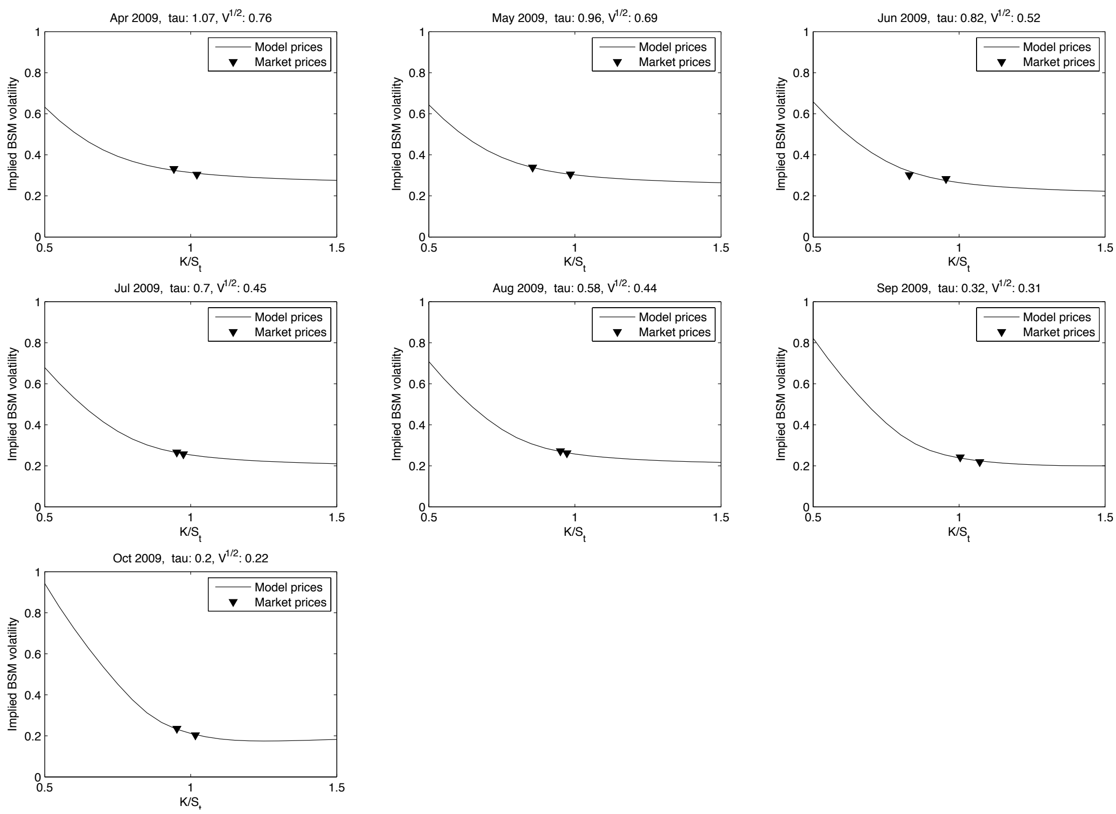

Figure 3.4: Plots of model and market prices of options on the OMXS30 price index expressed as implied volatilites.

Chapter 4

Simulation study

Having defined the model framework and the accompanying calibration procedure we now turn the attention to simulating long-term trajectories. In order to simulate these trajectories from the continuous-time model we define our discretization scheme of choice. Since we simulate long-term trajectories, i.e. the evaluation window exceeds the tenor of all available instruments, we need to introduce a set of investment strategies. Along with a reinvesting stock index strategy we define a bond strategy and a total of seven investment strategies involving structured products. The empirical study is then divided into two sections, where the first is a quantitative analysis of the terminal performance of the investment universe, whereas the latter is a scenario-based evaluation. These scenarios are chosen under the authours' discretion with the intent to capture some important features of these systematic investment vehicles.

4.1 Monte Carlo implementation

The simulation of the square-root processes describing the interest rate and the variance are well studied in the academic literature. It is a well-known fact that an Euler-Maruyama scheme generates negative values of the square-root processes with positive probability even when the Feller conditions are satisfied, i.e. when $2\kappa\theta > \sigma_V^2$ and $2\beta\alpha > \sigma_r^2$.[1] Hence, we follow the methodolgy of Goltz, Martellini and Simsek (2005) and apply a Milstein scheme that produces a smaller discretization error. Formally put, we use the discrete approximations $\widetilde{S}_t$, $\widetilde{V}_t$ and $\widetilde{r}_t$ as described in Equation 4.1.

$$
\begin{aligned}
\widetilde{S}_{t+\Delta} &= \widetilde{S}_t + \widetilde{S}_t\left(\widetilde{r}_t - d + \frac{1}{2}\widetilde{V}_t + \mu_0\right)\Delta + \widetilde{S}_t\sqrt{\widetilde{V}_t}Z^S_{t+\Delta}\sqrt{\Delta} + \widetilde{S}_t\left(J_{t+\Delta}-1\right)Q_{t+\Delta} + \widetilde{S}_t\frac{1}{2}\widetilde{V}_t\Delta\left({Z^S_{t+\Delta}}^2 - 1\right)\\
\widetilde{V}_{t+\Delta} &= \widetilde{V}_t + \kappa\left(\gamma - \widetilde{V}_t\right)\Delta + \sigma_V\sqrt{\widetilde{V}_t}Z^V_{t+\Delta}\sqrt{\Delta} + \frac{1}{4}\sigma_V^2\widetilde{V}_t\Delta\left({Z^V_{t+\Delta}}^2 - 1\right)\\
\widetilde{r}_{t+\Delta} &= \widetilde{r}_t + \beta\left(\alpha - \widetilde{r}_t\right)\Delta + \sigma_r\sqrt{\widetilde{r}_t}Z^r_{t+\Delta}\sqrt{\Delta} + \frac{1}{4}\sigma_r^2\widetilde{r}_t\Delta\left({Z^r_{t+\Delta}}^2 - 1\right),
\end{aligned}
\tag{4.1}
$$

where Z_t^S, Z_t^V, Z_t^r, Q_t and J_t are i.i.d. processes, $Q_t \backsim \mathrm{Po}\left(\Delta\lambda\right)$, $\ln\left(J_t\right) \sim \mathrm{N}\left(\mu_J, \sigma_J^2\right)$ and Z_t^S, Z_t^V, Z_t^r are $\mathrm{N}\left(0,1\right)$ variables with $\mathrm{Corr}\left(Z_t^S, Z_t^V\right) = \rho$. Furthermore, r_0 and V_0 are simulated from the unconditional distributions of r_t and V_t, i.e. the Gamma distribution, in an attempt to remove the bias of always starting in the same state.

In reality, the discrete approximation $\widetilde{S}_t$ involves a rather crude approximation of an Itô integral which reduces the number of simulations that needs to be performed per trajectory, but increases the numerical error. For the true multi-dimensional Milstein scheme see e.g. Kahl and Jäckel (2006). In our model, the numerical error appears to be extremely small compared to the stochastic error. This can be seen in Appendix B, where we use Monte Carlo simulations to verify the moment conditions derived in Section 2.2. As every trajectory requires the computation of a large number of option prices, the simulation progress is indeed very time-consuming. We hence settle for a total of 5000 trajectories.

[1] See e.g. Andersen (2007)

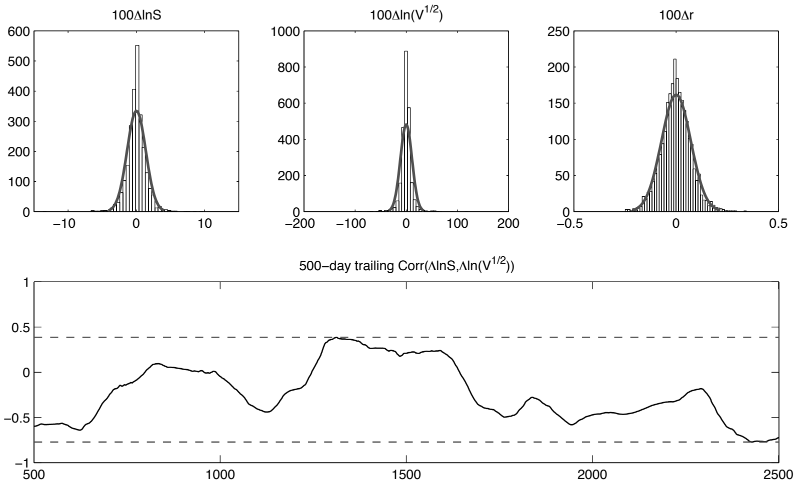

Figure 4.1: Upper: Histograms of $\Delta \ln S_t$, $\Delta \ln \sqrt{V_t}$ and Δr_t, respectively, from a randomly chosen trajectory over the 10-year evaluation window. Lower: Plot of the trailing 500-day correlation of $\Delta \ln S_t$ and $\Delta \ln \sqrt{V_t}$ of the same trajectory.

4.2 The investment universe

We restrict our investment universe to comprise

- an equity index investment vehicle,
- a bond strategy, and
- a class of portfolios of structured products.

The equity index investment vehicle is simply a reinvesting equivalent of the dividend-paying price index, often referred to as a gross index. To better explain the bond strategy and the class of portfolios of structured products we must introduce the concept of a roll.

A roll is a strategy dealing with securities that have a finite time to maturity or tenor, i.e. securities that eventually will mature, which in our model framework applies to the zero-coupon bonds (ZCBs) and the stock options. The roll is simply a strategy in which the investor purchases, say, a ZCB at time $t = 0$ with a tenor of 3 years. At time $t = 3$, i.e. when this ZCB matures, the proceeds are invested in a new, identical security. This procedure can be repeated infinitely many times. There are at least two important features to be mentioned about a roll strategy, namely that the investor has created a buy-and-hold-type of strategy that can be considered to have an infinite tenor and that the terminal payoff from such a strategy is unknown at time $t = 0$, as long as the investment horizon exceeds the tenor of the underlying security. In the case of the roll of ZCBs, the terminal payoff of a $(n \times \tau)$-year investment will be

$$\frac{100^{n+1}}{B_0^\tau \cdot B_\tau^\tau \cdot \ldots \cdot B_{(n-2)\tau}^\tau \cdot B_{(n-1)\tau}^\tau},$$

where B_t^τ denotes the price of a ZCB with tenor τ, issued at time t, whereas the 100 in the numerator refers to an initial investment of SEK 100.

4.2.1 Structured products

We now turn the attention to the class of structured products or equity-linked notes. The structured products we deal with are constructed so as to give the investor a known degree of capital guarantee while giving upside exposure to the equity market. This is done by combining the safety of a ZCB with the opportunities and risks associated with a leveraged instrument, namely a plain vanilla at-the-money call option.[2] Central in the discussion of such structured products are the concepts of participation rate and whether the product is issued at or above par. The participation rate can be interpreted as the share of nominal amount that is spent on options. A structured product of nominal amount SEK 100, issued at par, will have participation rate

$$PR = (100 - B_0)/C_0,$$

where PR is the participation rate while B_0 and C_0 are the prices of a ZCB and a plain vanilla, at-the-money option, respectively, at the time of issue (here $t = 0$).[3] Further, we introduce a margin and issue the structured product above par, that is we allow for the investor to invest SEK $(1 + \pi)100$ in the product while maintaining a capital guarantee equal to the nominal amount, here SEK 100. The prevailing participation rate is

$$PR = ((1 + \pi)100 - B_0)/C_0,$$

where $\pi > 0$ is the share of nominal amount above par. The larger the share spent on options, the lower the prevailing capital guarantee (as share of nominal amount). Thus, structured products issued above par would typically attract the more risk-prone investor. This terminology means that, while the nominal amount is guaranteed to be repaid at maturity, the invested capital is not. Rather, the prevailing capital guarantee, as share of invested capital, is reduced to $1/(1 + \pi)$. Figure 4.2 illustrates how the participation rate varies in π, r_t and V_t. Exposure-wise this means that the payoff at maturity (that is, at time $t = T$) of these structured productcs with a nominal amount of SEK 100 can be formulated as

$$100 + 100 \times PR \times \max\left(\frac{S_T - S_0}{S_0}, 0\right),$$

where S_t is the price of the underlying asset at time t and PR as specified above.

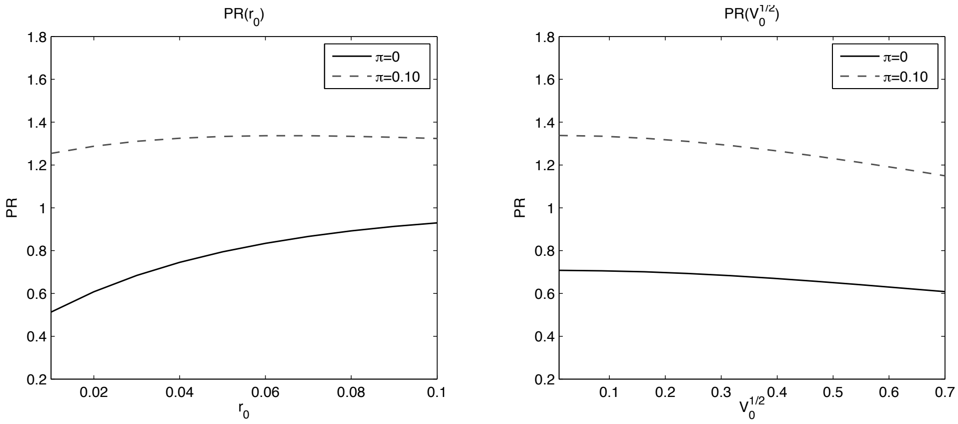

Figure 4.2: Prevailing participation rates for two distinct products specifications, namely 3-year structured products issued at and above par, i.e. with $\pi = 0$ and $\pi = 0.10$, respectively. Left: Participation rate as a function of the short rate with volatility equal to its long-term mean, i.e. $\sqrt{V_0} = 0.226$. Right: Participation rate as a function of the volatility level with the short rate equal to its long-term mean, i.e. $r_0 = 0.0316$.

[2]Most option constracts embedded in structured retail products are of more exotic art, e.g. containing Asian tails and/or Quanto features. This may render in cheaper as well as more expensive options and, consequently, higher as well as lower participation rates, although the main objective in the equity space is usually to eliminate unwanted risk exposures and increase the degree of participation.

[3]Note that the terminology at-the-money refers to the exercise price at the time of issue, i.e. that $K = S_0$.

4.2.2 Constructing the portfolios of structured products

The rolling procedure applies to structured products in the exact same fashion as for bonds. A roll of a structured product is thus an equity-linked buy-and-hold-type investment with a capital guarantee. This capital guarantee is, as we have seen, declining in π. Although the capital guarantee allows for a roll to lock in gains during the evaluation window, the roll only inherits the capital guarantee when the investment horizon is a multiple of the tenor of the structured products.

Further, in this study we assume a yearly issue frequency of these structured products of 10 issues per year. That is, every 25th business day, four different types of structured products are issued, varying in $\tau = [3, 4]$ and $\pi = [0, 0.10]$. Consequently, at any given point in time there is a total of 140 strucutred products on the market and, hence, 140 unique rolls.[4] In order for all unique rolls to be available for investment from time $t = 0$ we need to construct a pre-evaluation window, i.e. we need to evaluate the price paths of these rolls during a time period preceding the investment horizon. The pre-evaluation window length in years is $\tau - f^{-1}$, where τ is the tenor of the structured products and f denotes the issue frequency, in number of issues per year.

From these 140 rolls one could construct a structured products portfolio consisting of all 140 rolls or any subset thereof. However, to keep the investment universe comprehensible we define seven subsets and a specific weighting scheme and consider the resulting seven portfolios as the only tradeable strategies consisting of structured products. The portfolio constituents are displayed in Table 4.1 while the chosen weighting scheme is defined in Section 4.2.3.

	Tenor		At/above par		No of rolls	Pre-eval. window
	$\tau = 3$	$\tau = 4$	$\pi = 0$	$\pi = 0.10$		
Single roll	X	-	X	-	1	0 years
Portfolio 1	X	-	X	-	30	2.9 years
Portfolio 2	X	-	-	X	30	2.9 years
Portfolio 3	X	-	X	X	60	2.9 years
Portfolio 4	X	X	X	-	70	3.9 years
Portfolio 5	X	X	-	X	70	3.9 years
Portfolio 6	X	X	X	X	140	3.9 years

Table 4.1: Specification of portfolio constituents in Single roll and Portfolios 1 through 6.

[4]The rationale behind the total number of unique rolls is that for a portfolio containing structured products with a 3-year tenor only, then $\tau \times f = 3 \times 10 = 30$ rolls are needed to reach an euquillibrium state. Using the same analogy we conclude that 40 rolls are needed for a portfolio of 4-year structured products to reach an equillibrium state. Since there are two distinct product types, that is a structure products can be issued at ($\pi = 0$) or above ($\pi = 0.10$) par and, consequently we need $2 \times (30 + 40) = 140$ unique rolls.

4.2.3 Portfolio weighting scheme

Given the portfolio constituents in Table 4.1 we define a weighting scheme to explicitly determine the portfolio compositions. In order to understand the chosen weighting scheme, one should be aware that for a given trajectory the start time t_0 of the evaluation window will impact the future performance of any roll. This stems from the fact that the structured products constituting the roll are sold at $t = t_0$, $t_0 + \tau$, $t_0 + 2\tau$, ..., and, hence, the state variables $V_{t_0+k\tau}$ and $r_{t_0+k\tau}$ determine the participation rates while, since the options are sold at-the-money, $S_{t_0+k\tau}$ determines the strike prices. In Section 4.3 any bias stemming from this dependence is handled by letting the starting values V_0 and r_0 be random variables themselves, but in a discretionary setting such as in Section 4.4 this is an important feature of the investment alternatives.

Weighting of the rolls in the abovementioned portfolios may be done either through a capital-weighted or an absolute scheme. A capital-weighted scheme would imply that in-the-money products are underweighted compared to out-of-the money products. Thus, such weighting scheme distorts the inherent properties of the structured products as bonds would be assigned higher weights than would the average structured product. Furthermore, the average tenor of the portfolios would be lower than should one use an absolute scheme.

Instead, we suggest using an absolute weighting scheme. A naïve absolute weighting of any Portfolio 1 through 6 consisting of n rolls would be to invest equal amounts in each roll, i.e. the weight α_t^i assigned to each roll i would be defined as

$$\alpha_t^i = \frac{V_t^P}{\sum_{j=1}^n V_t^j},$$

where V_t^P is the portfolio value and V_t^i is the value of roll i. Note that, according to this scheme, the weights α_t^i are constant in time and hence require no rebalancing. This scheme would however be highly path dependent as the previous performance of each roll would determine the weight assigned to each product. Hence, we suggest an equal capital protection scheme which assigns equal absolute weights to the products rather than the rolls, i.e. the weights are defined as

$$\alpha_t^i = \frac{V_t^P \sum_{j=1}^n nbonds_t^j}{nbonds_t^i \sum_{j=1}^n V_t^j},$$

where $nbonds^i$ is the number of bonds held in roll i at time t. Since the number of bonds held in each roll only changes when a product matures this scheme requires rebalancing ten times per year.

4.3 The quantitative approach

Performance analysis in finance generally revolves around deducing a relationship between the risks and returns associated with a specific financial product or instrument. Several popular ways of measuring financial performance are based on the mean and variance/standard deviation of historical returns (e.g. Markowitz' mean-variance frontiers and the Sharpe ratio).

When dealing with non-Gaussian return distributions several questions can be raised regarding how to quantify risks. Return distributions can show non-Gaussian properties by exhibiting skewness, excess kurtosis or both. In the case of asymmetric return distribution it is imperative to separate the notions of risk and opportunity, as they are interchangeable only in a symmetric setting. For instance, a large positive skew with extreme return shocks in the positive tail would typically show a large standard deviation, although such a deviation from normality should by no means be punished by being accompanied with a high level of risk. Merely, this asymmetry represents a high level of opportunity (not unlike a lottery ticket), as long as the investor prefers more money to less. Another example would be option contracts as they show convex relation to the underlying asset rendering in skewed return distributions even in a Gaussian framework such as the BSM model. Hence, we prefer quantile-based measures to those incorporating an enire distribution function.

Regardging excess kurtosis or so called fat tails, the rationale is that a fat-tailed distribution would more frequently exhibit extreme outcomes as well as outcomes centered around the mean than would a Gaussian distribution with the same mean and variance. Consequently, a parametric approach to measuring tail risk would fail in such a setting. On back of these arguments, along with our simulated data showing both skewness and excess kurtosis, we rely on a non-parametric, quantile-based risk measures.

Before stipulating our risk measures of choice we need a formal definition of how to deal with the terminal returns. In an attempt to reduce the extreme behavior of the total returns we choose to evaluate the geometric means of the 5000 terminal, ten-year returns. That is, given a total return $R_{0,10}$, the geometric mean, or Compounded Annual Growth Rate (CAGR), is defined as

$$CAGR_{0,10} = (1 + R_{0,10})^{1/10} - 1.$$

In order to specify the risk measures we let $L = -CAGR_{0,10}$, where L is to be interpreted as the geometric average annual loss. Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) are then defined

$$\mathrm{VaR}_\alpha(L) := \inf\{l \in \mathbb{R} : \mathrm{P}(L > l) \leq 1 - \alpha\},$$

and

$$\mathrm{CVaR}_\alpha(L) := \mathrm{E}[L \mid L \geq \mathrm{VaR}_\alpha(L)].$$

We note that the higher the reading of VaR or CVaR, the higher the level of risk associated with the instrument. Being a simulation study we go about the calculation of these risk measures in a purely empirical way.[5] Figure 4.3 shows how CVaR varies in the threshold parameter α for three of the nine different investment strategies.

4.3.1 Introducing the Single roll to a classical stock-bond mix

We start out by studying the simplest strategy involving structured products, namely the Single roll. A brief inspection of Table 4.2 reveals that the bond strategy offers a safe haven as it shows low variance and negative readings of both VaR and CVaR. The stock strategy allows for substantially higher returns than do bonds, yet is associated with higher levels of risk judging by all available measures. Finally, as one would expect, the Single roll offers a compromise between the two. Further, judging by the scatter plots shown in Figure 4.4 we find no significant correlation pattern between neither bonds and stocks nor bonds and Single roll. Consequently, this leaves us to conclude that the main driver for Single roll returns is stock performance, as the two show a clear convex return relationship.

[5] The empirical VaR on threshold level $\alpha = 0.01$ is defined as the 50th largest loss, while CVaR, accordingly, is the arithmetic average of the 50 largest losses.

As a sanity run-through of the terminal returns we conclude that stocks outperformed bonds four times as often as the other way around. Specifically, the former occured approximately 81% of the times, and when doing so, Single roll was the second best strategy 86% of the times. Similarly, among the 19% of the simulated trajectories where bonds oputperformed stocks we see that Single roll was the second best strategy 57% of the times. Further, among the times where bonds outperform stocks we notice a slight skew towards Single roll outperformning both assets more often than underperforming both of them. Unconditionally, that is regardless of stocks outperforming bonds or the other way around, we note a fairly symmetric distribution with Single roll being the best alternative 9%, second best 81% and the worst 10%.

To quantitatively deduce and illustrate whether or not to include these instruments in our portfolio we draw efficient frontiers. That is, we find portfolio weights so as to maximize the expected return for a given level of expected risk, measured as variance and CVaR, respectively.[6] Geometric means of the simulated returns are used as proxies for future expectations of the one-period model. From Figure 4.5 we see that although a small portion is attributed to the Single roll there is no significant improvement in the reachable mean-variance optimal portfolios. On the contrary, when risks are measured as Conditional Value-at-Risk we note improvements on both ends of the risk-return spectrum as the Single roll crowds out stocks in the lower risk-return region, whereas in the upper risk-return region bonds are crowded out by the Single roll. This crowding out is showed in Figure 4.6.

[6] Or, equivalently, weights are chosen so as to minimize the expected risk for a given expected return.

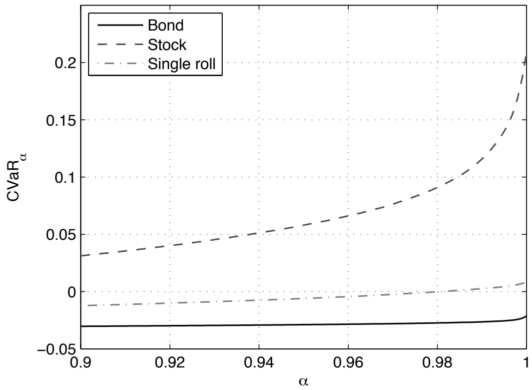

Figure 4.3: Conditional Value-at-Risk (CVaR) varying in the threshold parameter α for the bond, stock and Single roll strategies.

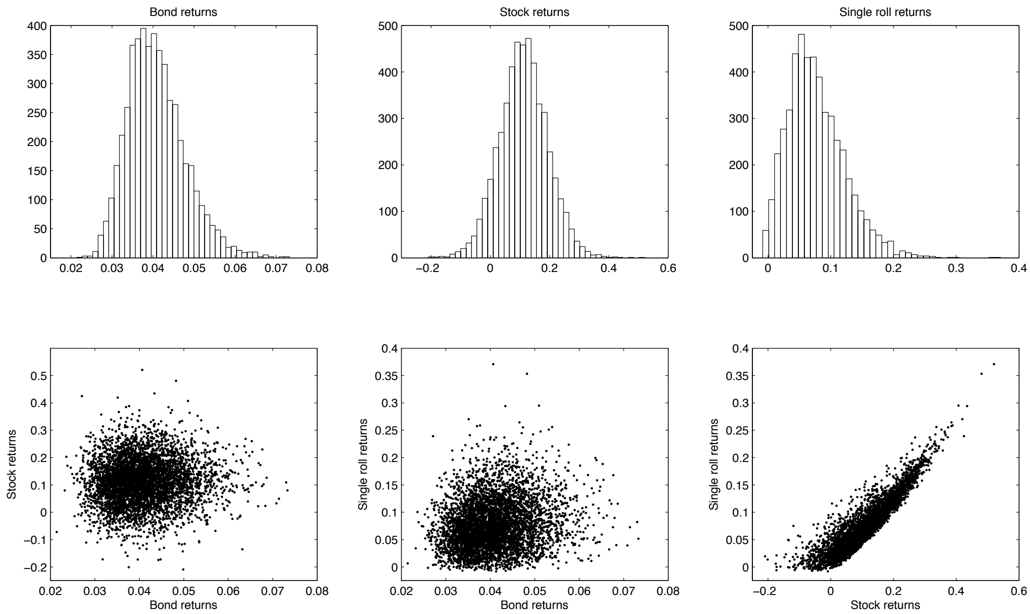

Figure 4.4: Histograms and scatter plots of the stock, bond and Single roll strategies. Returns on CAGR basis.

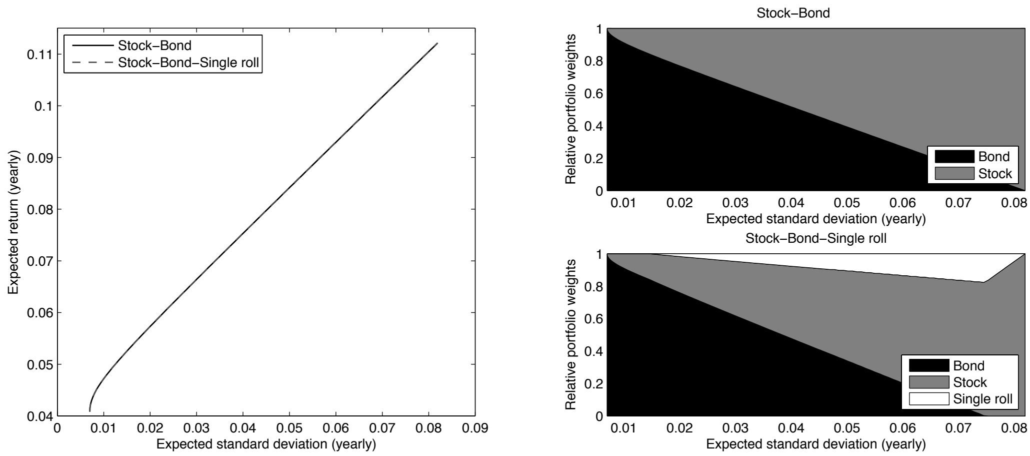

Figure 4.5: Efficient frontiers based on simulated data. Returns on CAGR basis and risk measured as standard deviation.

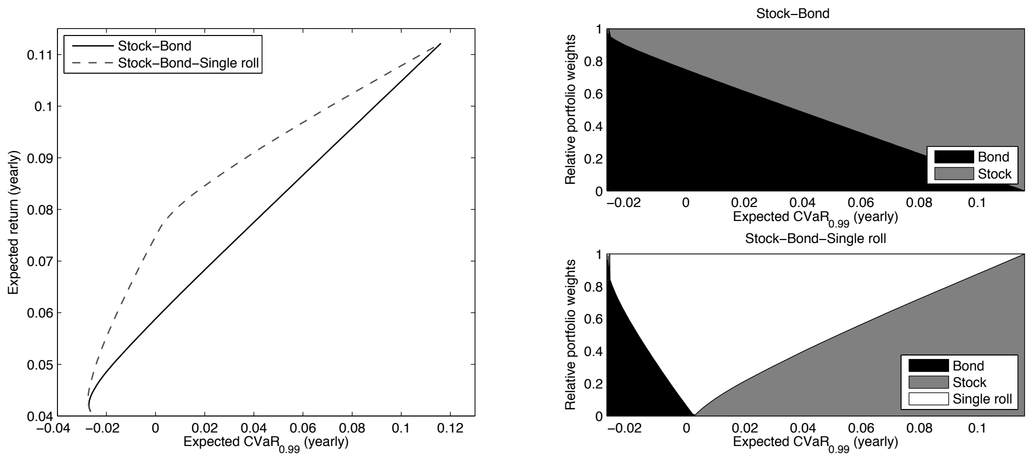

Figure 4.6: Efficient frontiers based on simulated data. Returns on CAGR basis and risk measured as Conditional Value-at-Risk.

4.3.2 Introducing path dependence via a portfolio of multiple rolls

Having pointed out some of the advantages and drawbacks of structured products examplified by the Sinlge roll we now look at how this Single roll compares to the remaining strategies in the class of systematic investment vehicles involving structured products. As explained in Section 4.2.2, the Single roll can be viewed as a subset of Portfolio 1. Consequently, we focus on the differences between the two.

We know that as the Single roll contains one roll only there is no need for a pre-evaluation window. Hence, we say that Single roll is state dependent in the sense that the prevailing participation rate of the first structured product (issued at time $t = 0$) is dependent upon the state variables V_0 and r_0. Moreover, Portfolio 1 is said to be path dependent as not only the state variables but also the price paths of the underlying securities within the pre-evaluation window impact the composition of the portfolio at time $t = 0$. As a first comparison, we study the static measures as shown in Table 4.2 to find that not a whole lot can be said about the difference between the two strategies.

In a further attempt to disentangle the effects of state dependence from those of path dependence we study the difference in performance between Portfolio 1 and the Single roll, i.e. how much Portfolio 1 outperforms Single roll over the ten year long evaluation window, on CAGR basis. This difference is shown in Figure 4.7. We find no significant correlation between this difference and stock performance, neither that prior to the evaluation window nor that of the actual evaluation window. The fact that we cannot find any distinct relation between the performance of the underlying leads us to think that the difference can be derived to what we refer to as intra-evaluation window performance. In other words, the difference stems from where the local highs and lows are within the evaluation window. We call this timing risk and the matter is discussed further in Section 4.4. However, we choose to present the two most extreme outcomes of this difference in Figure 4.8.

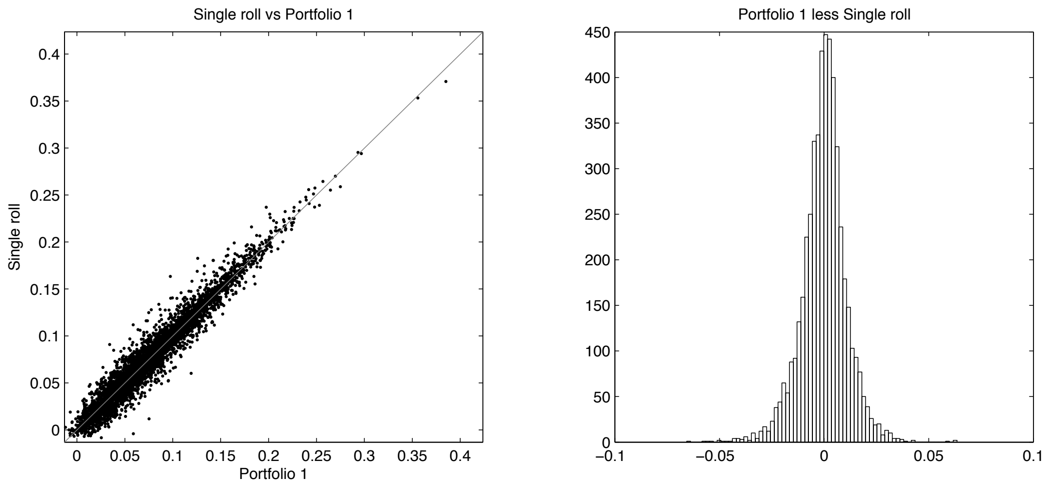

Figure 4.7: Left: Scatter plot of Single roll returns vs Portfolio 1 returns. Right: Histogram of the difference between the two.

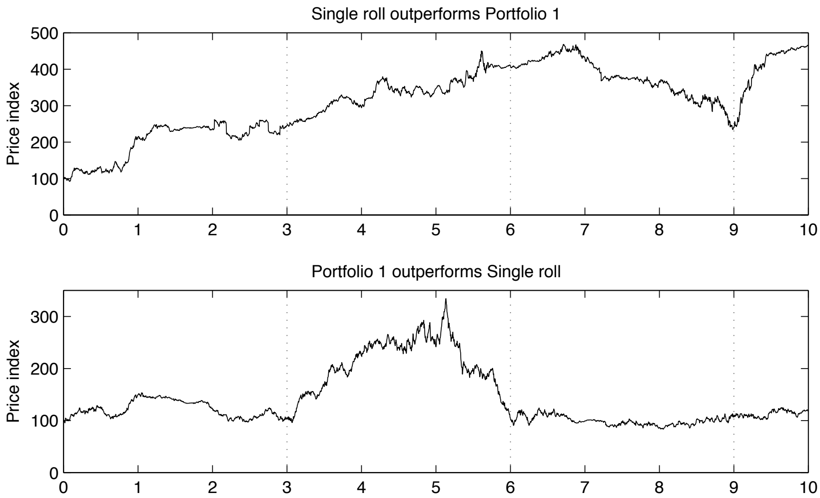

Figure 4.8: Plots of the evolution of the price index for the two scenarios producing the largest differences between Single roll and Portfolio 1. The dashed lines indicate where the structured products in the Single roll matures and new ones are issued. Upper: Total returns of 153% and 353% for Portfolio 1 and Single roll, respectively. Lower: Total returns of 107% and 12% for Portfolio 1 and Single roll, respectively.

4.3.3 Diversifying in tenor and capital guarantee via Portfolios 2-6

Now that we have derived the effects of introducing multiple rolls we look at how we can alter the performance palette by diversifying in tenor and issuing structured products above par. Figure 4.9 shows the relationship between mean return and CVaR, where we conclude from the clustering of portfolios that diversification in tenor seems to have little effect. In other words, Portfolios 4 through 6 provide little extra information on a risk-return basis.

On a total return basis we see that, although there is no capital guarantee, CVaR is still very low for strategies involving only structured products issued at par. Specifically, we find total return CVaRs of 2.5% and 0.6% for the Single roll and Portfolio 1, respectively. When introducing structured products issued above par we find higher risk levels and, by the same token, they offer more attractive returns. As an example, we note that as compared to stocks with a mean total return of 268% and corresponding CVaR reading of 69.4% the mixed portfolio Portfolio 3 shows a mean total return of 168% and CVaR of 14.6%. Frequency-wise, we conclude that roughly 3 out of 4 times Portfolio 2 is superior to Portfolio 1, i.e. the extra money spent on options is worthwhile 75% of the time. The scatter plots in Figure 4.10 reiterate this as the main difference between Portfolios 1 and 2 is the slope of the positive tail rather than the worst case scenarios being less extreme for the former portfolio.

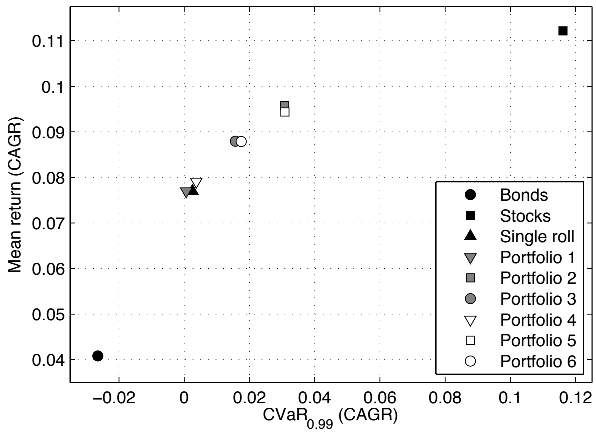

Figure 4.9: Plot of the prevailing risk-return relations of the nine investment strategies, as measured by simulated data. Expected return and and the corresponding $\text{CVaR}_{0.99}$ on CAGR basis.

	Mean	St. dev.	Skewness	Kurtosis	Min	Max	$\text{VaR}_{0.99}$	$\text{CVaR}_{0.99}$
Bonds	4.1%	0.7%	0.67	3.70	2.1%	7.3%	-2.8%	-2.6%
Stocks	11.2%	8.2%	0.06	3.45	-20.9%	52.1%	8.1%	11.6%
Single roll	7.7%	4.5%	0.84	4.21	-0.8%	37.1%	-0.0%	0.3%
Portfolio 1	7.7%	4.4%	0.90	4.48	-1.2%	38.5%	-0.5%	0.1%
Portfolio 2	9.6%	6.9%	0.75	3.88	-5.0%	52.0%	2.1%	3.1%
Portfolio 3	8.8%	5.8%	0.85	4.21	-3.1%	47.1%	0.9%	1.6%
Portfolio 4	7.9%	4.7%	0.89	4.39	-1.6%	40.5%	-0.2%	0.4%
Portfolio 5	9.4%	6.9%	0.76	3.85	-5.0%	51.6%	2.3%	3.1%
Portfolio 6	8.8%	6.0%	0.84	4.13	-3.3%	47.3%	1.1%	1.7%

Table 4.2: Summary statistics of the 5000 simulated trajectories. Returns and associated measures on CAGR basis.

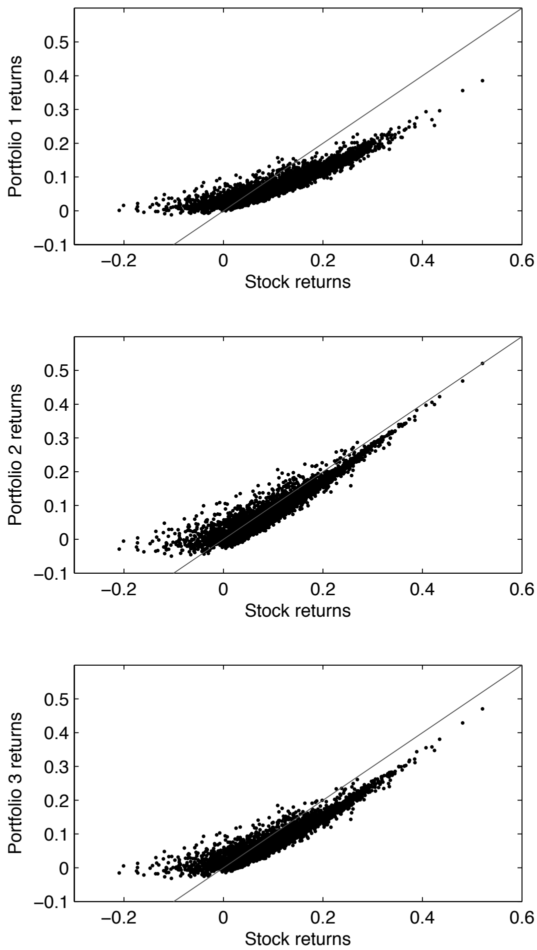

Figure 4.10: Scatter plots of Portfolios 1, 2 and 3 against a gross equity investment. The line indicates a strictly linear relationship.

4.4 Scenario analysis – the discretionary approach

The quantitative analysis in Section 4.3, albeit informative, does not cover a few issues that are best illustrated by looking in depth at single trajectories. The main concern with a purely quantitative approach is that it fails to disentangle market risk, which we define as the risk of chocks to market factors, from timing risk, which we define as the risk induced by uncertainty of the market state at the time of investment. Accordingly, we have chosen three trajectories of interest, namely the trajectories displayed in Figures 4.11, 4.13 and 4.15.

The first scenario in this analysis exhibits a distinct bubble formation which illustrates some of the issues with market timing. In this environment there is plenty of opportunities for an active portfolio manager to out- or under perform the market while a passive investment in stocks would generate a mere 29% ten year reutnr including dividends.

Similarly, scenario 2 displays some rather extreme market highs and lows which, naturally, is beneficial for the portfolio of structured products. In this scenario, the single roll outperforms the weighted portfolios since the stock index returns the first three years as well as the three years following $t = 6$ are very favorable. Again, the starting time of the Single roll has a significant impact on the terminal return.

Scenario 3 describes a market where a large drawdown during the fourth year in stock index level is not recovered during the following six year period. During this scenario the Single roll strategy does not even perform at par with a pure bond roll and the stock index displays a negative (including dividends) 10 year return. Even though there are some local highs, including a very strong year 6, these are not enough for Portfolio 1 to perform at par with bonds. Portfolio 6 performs even worse due to the increased weight assigned to options.

4.4.1 The time-variability of market risk

Since the relative weighting of bonds and options in structured products are time-varying, so is the instantaneous market risk of a portfolio or single roll. In this study we introduce a quantile-based risk measure, inspired by VaR, to represent market risk. The percentile measure is the relative loss given an instantaneous and simultaneous relative decrease in stock index price, increase in interest rate and decrease in the variance process corresponding to the 0.1% most extreme days. Formally put, we define the risk measure as follows described in Equation 4.2.

$$risk_t = 1 - \frac{V^P\left(S_t\left(1+S^*\right), V_t\left(1+V^*\right), r_t\left(1+r^*\right), t\right)}{V^P\left(S_t, V_t, r_t, t\right)}, \tag{4.2}$$

where,

$$P\left(\frac{S_{t+\Delta}}{S_t} \leq 1 + S^*\right) = 0.001$$
$$P\left(\frac{V_{t+\Delta}}{V_t} \leq 1 + V^*\right) = 0.001$$
$$P\left(\frac{r_{t+\Delta}}{r_t} \geq 1 + r^*\right) = 0.001,$$

and our estimates using the simulated daily returns are $S^* = -0.064$, $V^* = -0.82$ and $r^* = 0.098$. As expected, the average level of market risk, in all three scenarios, is very similar between Portfolio 1 and the Single roll while the market risk is slightly more volatile for the single roll, see Figures 4.12, 4.14 and 4.16. Furthermore, the market risk in either strategy is highly dependent on the stock index performance. This is not very surprising as the relative weights assigned to options as well the Delta of the options are increasing in stock returns and hence, this is an inherent property of a principal protected note. Still, this property should be taken into consideration when investing. Firstly, a structured product at the time of issue could be considered a low-risk investment, a property that is quickly eliminated by a strong performance of the underlying. Secondly, as the stock market falls heavily, the value of option contracts embedded in the structured products deteriorates and hence there is little room for a recovery of the product. This can be seen in particular in Figure 4.16 where the risk of the Single roll is very close to zero during the sixth and the eighth years and does not react significantly to the variablity of the stock market during those years. It is interesting that this market risk variability is partially hedged by investing in a weighted portfolio, in particular the risk does not approach zero for the portfolios. Should an investor wish to reduce the volatility of market risk further, a dynamic hedging strategy involving stocks and/or bonds is probably required. Finally, Portfolio 6 does not differ noticeably in risk volatility from Portfolio 1 while the average level is higher as a result of the increased portfolio weights assigned to options.

4.4.2 Market timing and path dependence

Timing risk, which is particularly interesting due to the properties of structured products, is an important issue from an investor perspective that can also be considered a state dependent and hence time-varying risk. In particular, as shown in Section 4.2.3 the future performance of the Single roll is highly dependent on the chosen time of investment. The quantitative approach does account for this type of risk as the starting states are themselves random variables. However, the large number of trajectories makes it hard, if not impossible, to disentangle timing risk from market risk, and the very nature of this risk factor makes it very hard to quantify numerically.

Studying Figure 4.12, the advantages of structured products in Scenario 1 are clear. The weighted portfolios of structured products locks in the gains during the market rise of the first four years of the period and distinctly outperforms both stocks and bonds. The base strategy of a Single roll has similar characteristics but does not eliminate the timing risk to the same extent. In particular, the Single roll strategy with products maturing at $t = 3, 6, 9$ years does not benefit from the stock market highs during the fifth and sixth years. Clearly, starting the roll at different points in time generates different returns while the weighted portfolio, which is standardized ten times per year, is not as dependent on starting point.

In Scenario 2, we see an inverted pattern. The market highs at $t = 3, 9$ provides an impressive return of the Single roll. If we study the market risk as illustrated by the quantile risk measure we find that the level of market risk does not differ significantly between the Single roll and Portfolio 1, in fact the average risk only differs by 0.02%. Still, there is obviously a significant difference of the return profiles of these strategies and how their performance relate to the trajectories of underlying market factors.

The static properties of a Single roll allow the future portfolio weights, and hence performance, to be determined partially by the time of investment. Although this timing risk is not fully displayed in the quantitative section, we have in these scenarios shown that market timing presents a real risk that can be exploited by well informed investors or hedged via a weighted portfolio.

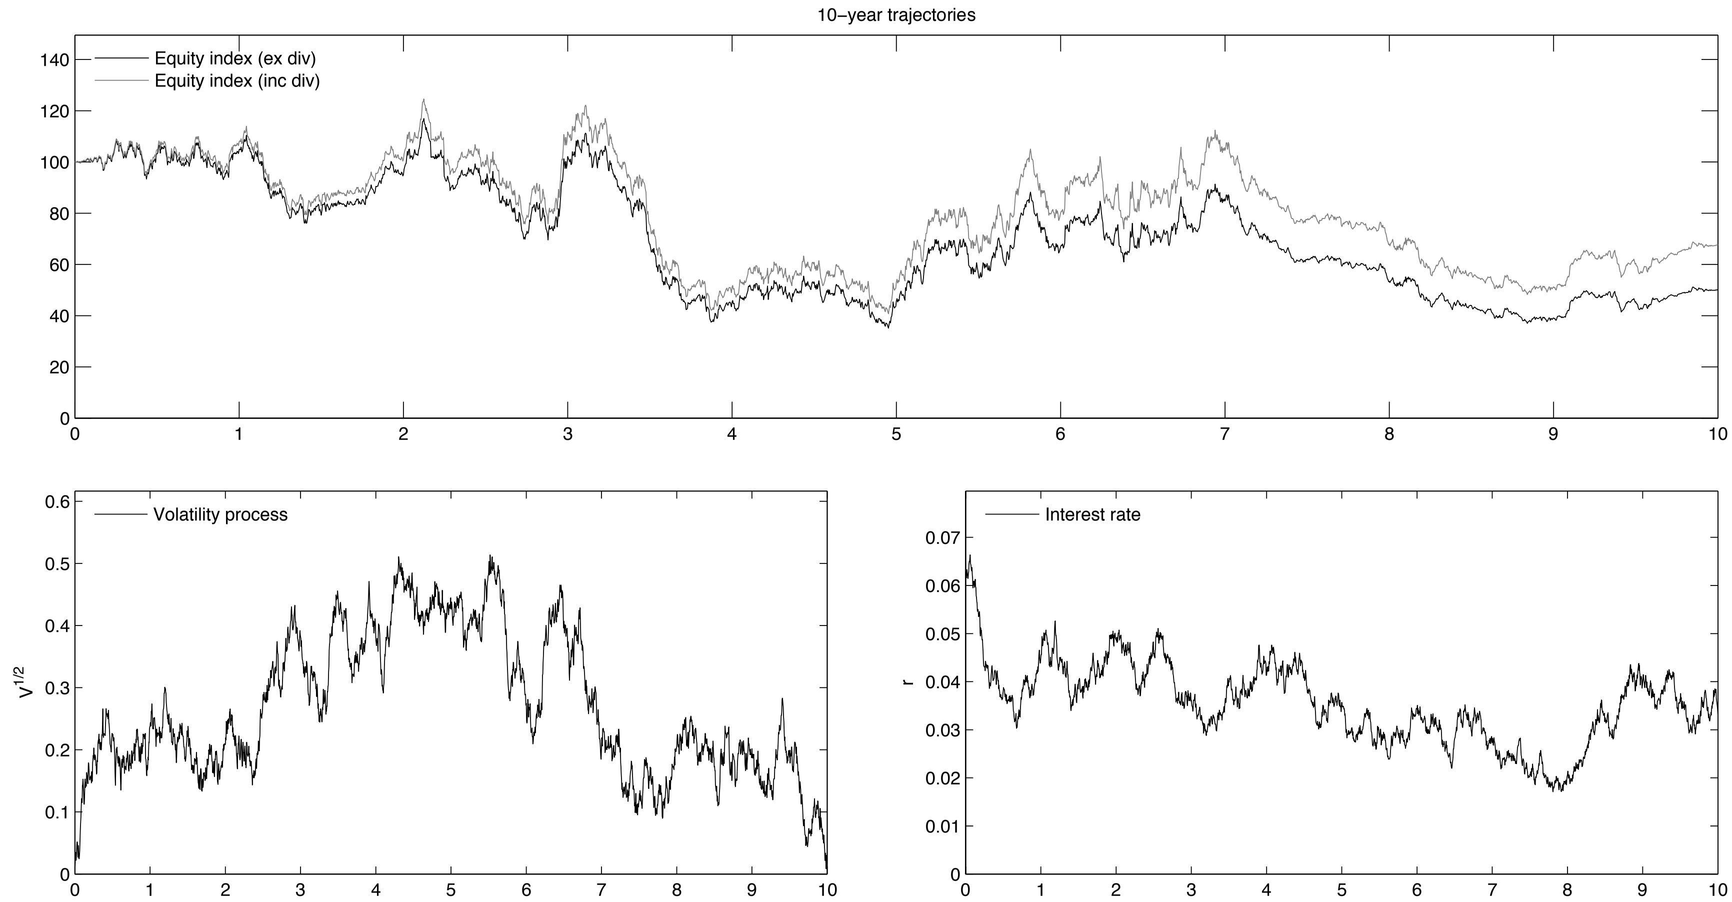

Figure 4.11: Scenario 1. Trajectories of the state variables.

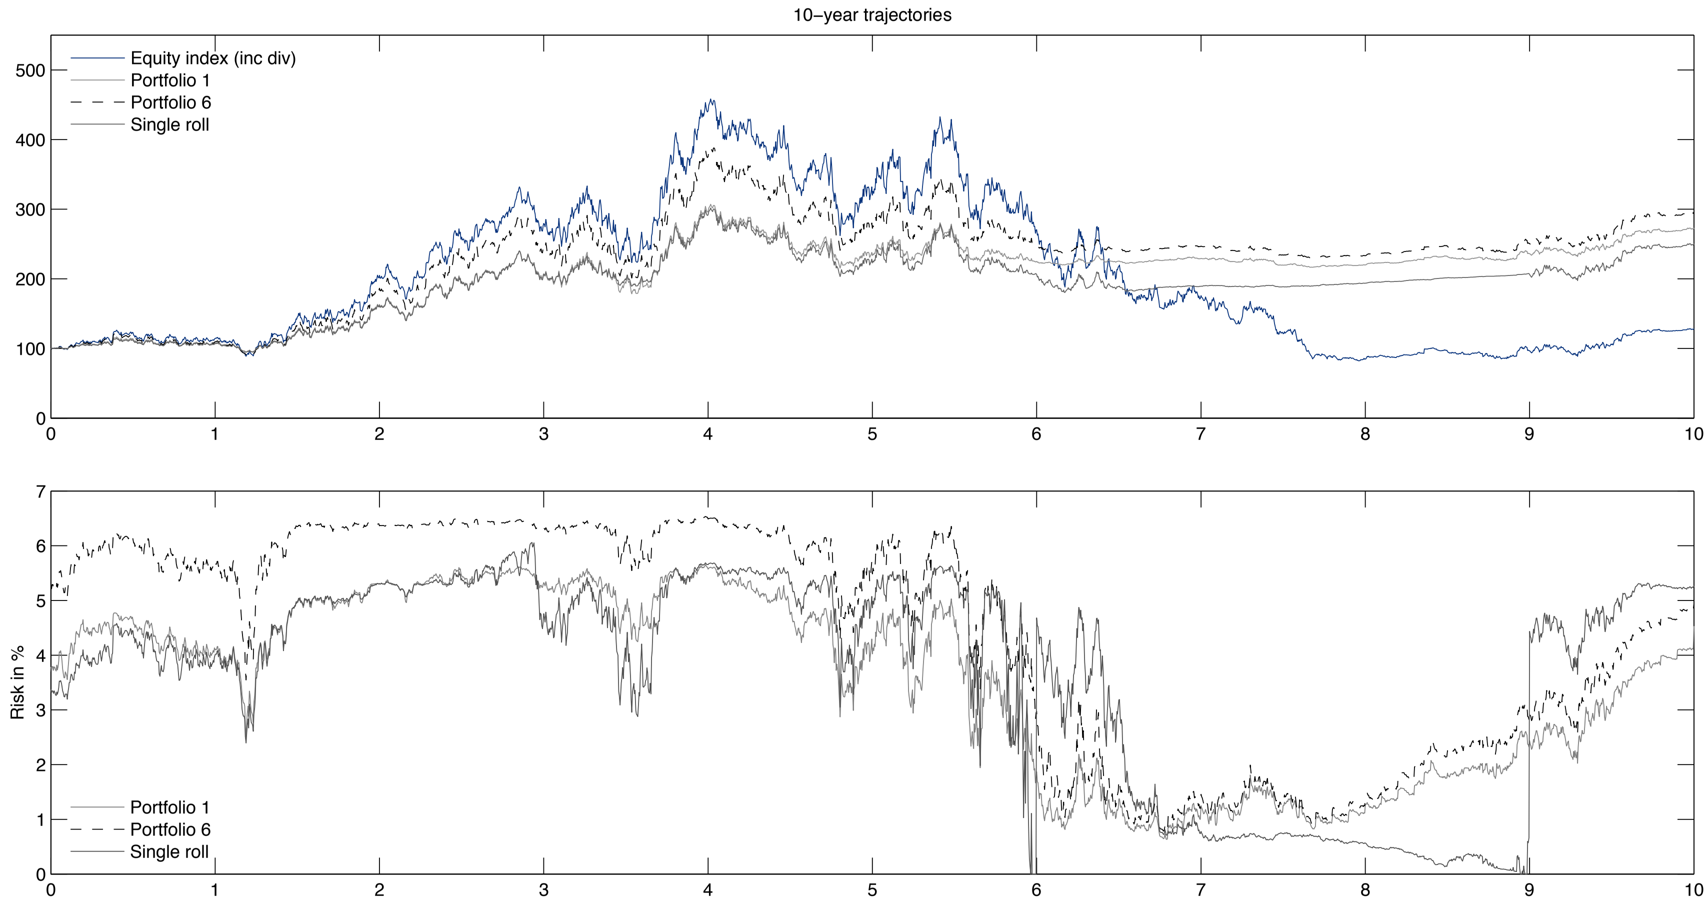

Figure 4.12: Scenario 1. Trajectories of the investment alternatives.

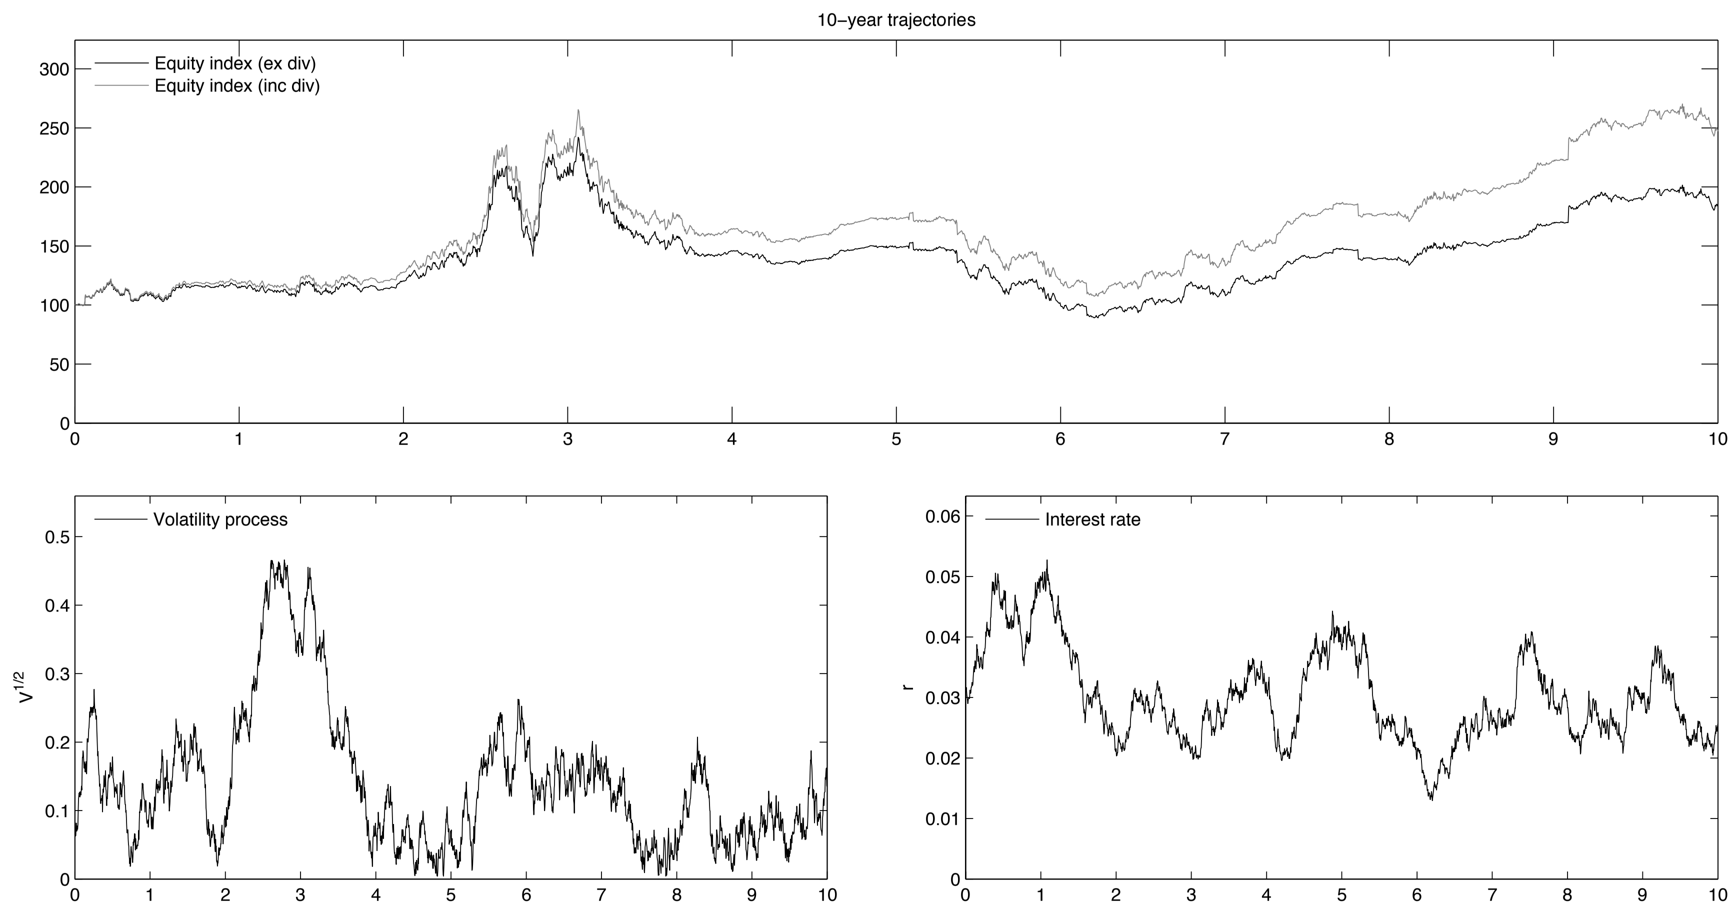

Figure 4.13: Scenario 2. Trajectories of the state variables.

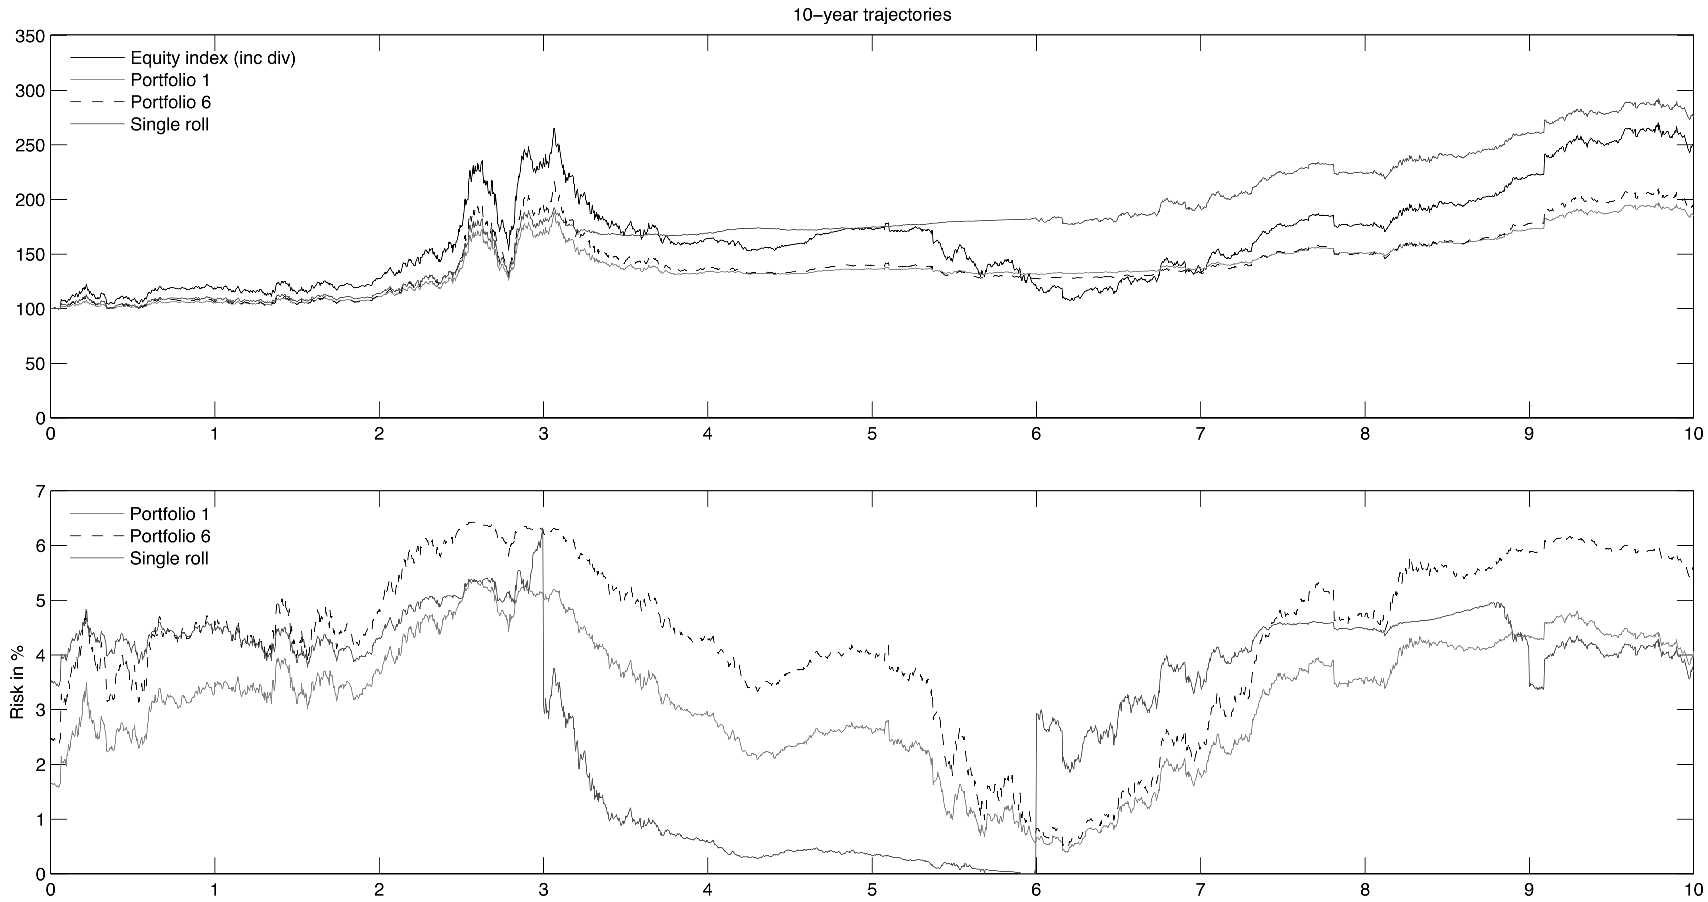

Figure 4.14: Scenario 2. Trajectories of the investment alternatives

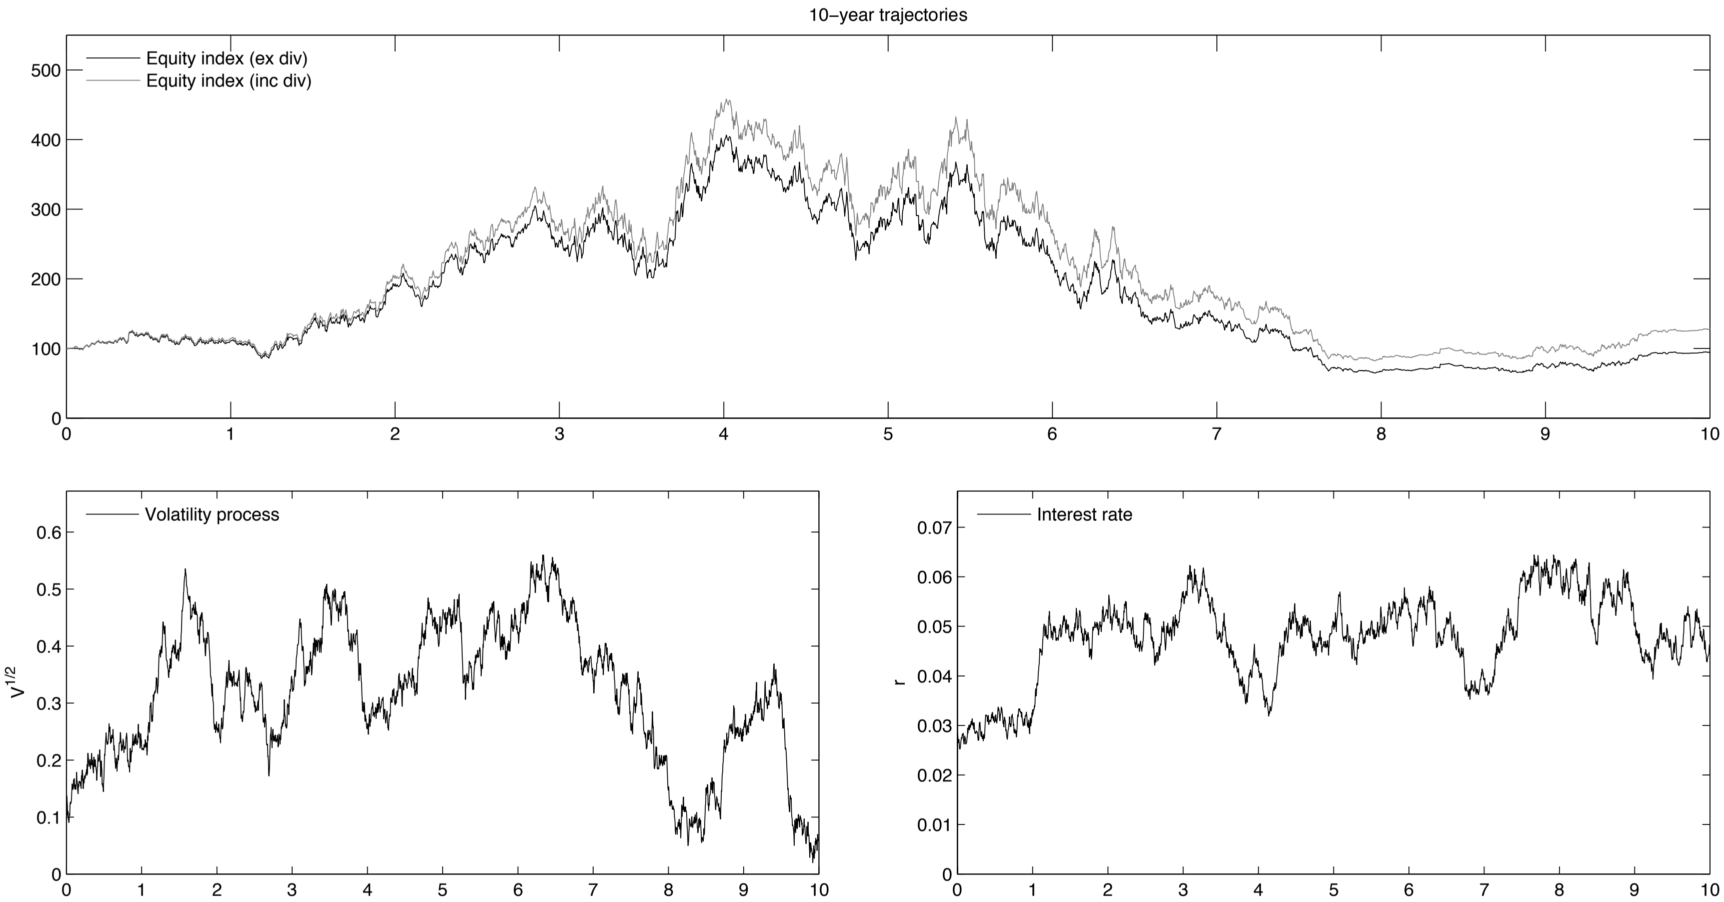

Figure 4.15: Scenario 3. Trajectories of the state variables.

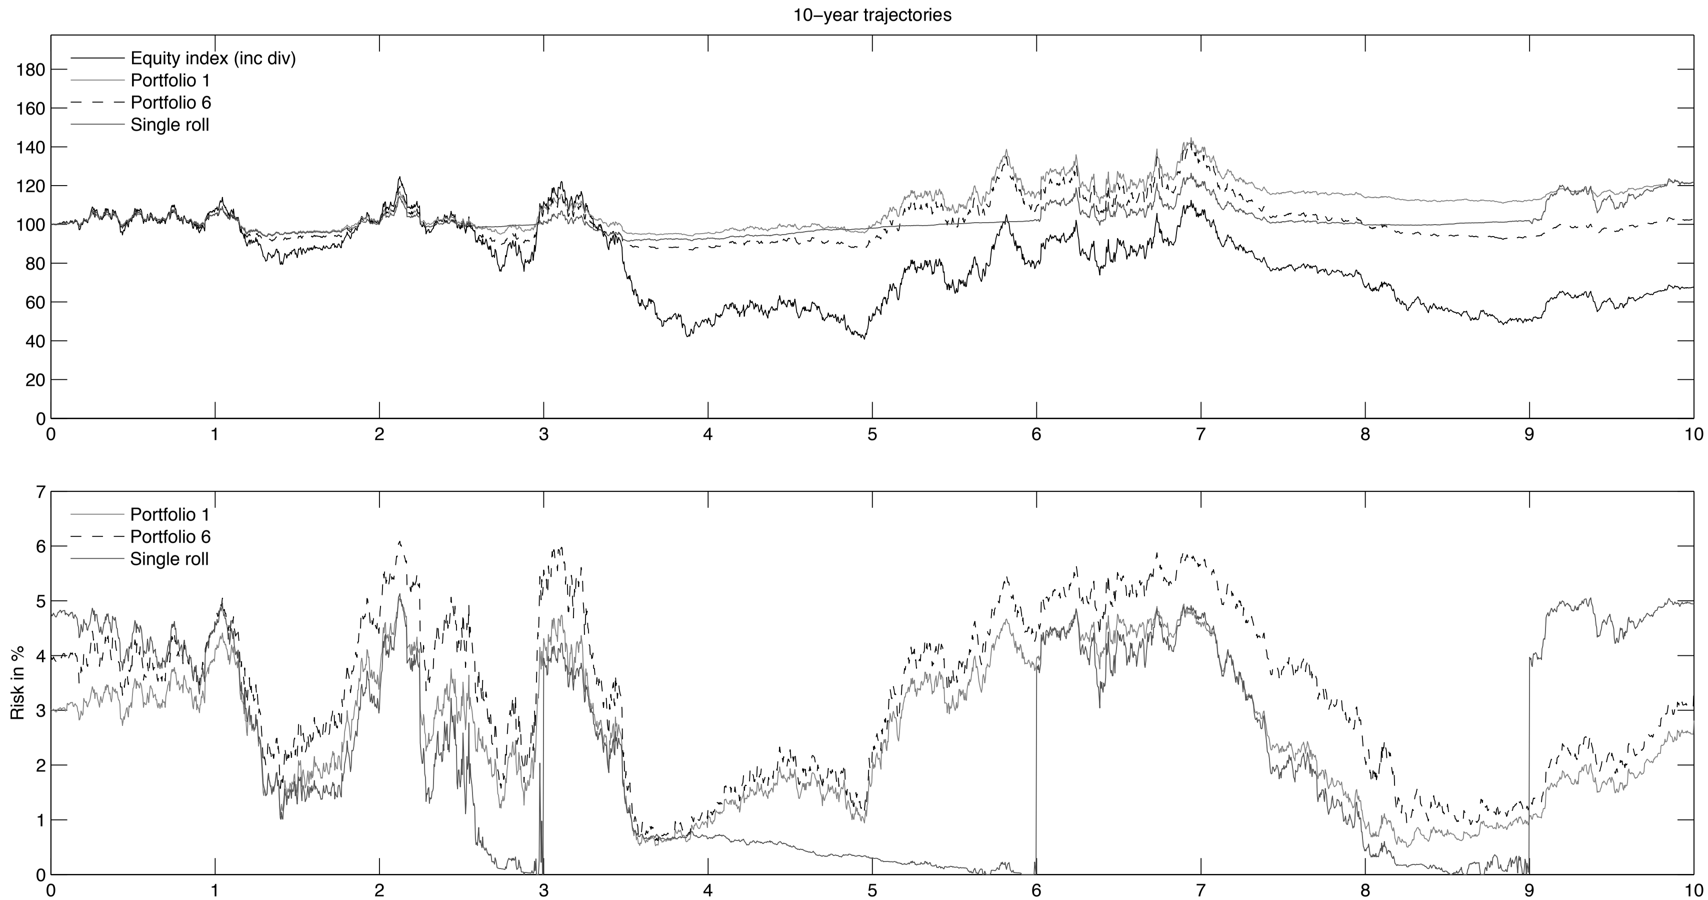

Figure 4.16: Scenario 3. Trajectories of the investment alternatives.

Chapter 5

Conclusions

This thesis uses a Monte Carlo approach to examine the properties of systematic investment vehicles consisting of structured products in a setting where the investment horizon exceeds the evaluation window. In such a setting it is not possible to fully benefit from the the capital guarantee and, hence, the products exhibit a slightly different risk profile. Although the capital guarantee is removed, the structured products returns still exhibit a distinct convexity relative to the underlying equity index. Further, among the trajectories where bonds outperform stocks we notice a slight skew towards Single roll outperforming both assets.

We employ a stochastic volatility, random jump model incorporating stochastic interest rates that proves to provide superior explanatory power over the Black-Scholes-Merton model in terms of historical model fit. A Milstein discretization scheme is implemented to simulate plausible future price scenarios and for the according trajectories, price elementary and derivative instruments, capturing both volatility skew with a varying term structure and a non-flat time-varying yield curve.

As simulated asset returns exhibit skewness as well as excess kurtosis we opt for a quantile-based risk measure and conclude that the introduction of structured products significantly improves the risk-return space available to investors. However, in a mean-variance setting structured products provide no improvement as the benefit of a positive skew is not considered. We take this as evidence of the inappropriateness of variance as a measure of risk rather than evidence against the existence of structured products.

Further, we show that holding multiple products helps in reducing what we refer to as timing risk. This finding is of particular interest as the timing risk may be a key reason for an investor to choose structured products in the first place. Timing risk is said to stem from intra-evaluation window performance, i.e. the risk (or opportunity) of products maturing and thereby reinvesting where the underlying asset, the equity index, exhibits a local extreme. On a more qualitative note, we see that the portfolios of structured products, i.e. multiple rolling investments, slightly reduce the time-variability of market risk as there is a smoothening effect on the relative weighting between bonds and options, respectively. When we diversify in tenor and capital guarantee we see that even though the portfolios involving structured products issued above par or with longer time to maturity increase leverage they show attractive risk measures as the feature of multiple rolls, in almost any market climate, lock in gains from one or more local highs rendering in a cumulative non-zero option payoff over a ten-year period. Studying simulated returns we find that a portfolio consisting of above-par products outperforms a portfolio of products issued at par three out of four times.

Fees are always present for any investor and they come in many forms and could well provide an interesting topic for further studies. Moreover, if a portfolio or indexed product like those described in this thesis where to be sold there would be no apparent reason to restrict the products to options on the same underlying, hence a future study could benefit from including multiple underlying securities, be it to represent different asset classes, regions etc. In particular if the underlying of the portfolio constituents are not perfectly correlated the risk profile of the portfolio would most likely prove to be less volatile which would provide a more transparent investment alternative.

Bibliography

[1] Andersen, L. (2007), *Efficient Simulation of the Heston Stochastic Volatility Model*, Working Paper, Bank of America.

[2] Andersen, T.G., L. Benzoni and J. Lund (2002), *An Empirical Investigation of Continuous-Time Equity Return Models*, vol. 57, pages 1239–1284.

[3] Andersen, T.G., L. Benzoni and J. Lund (2004), *Stochastic Volatility, Mean Drift and Jumps in the Short Rate Diffusion: Sources of Steepness, Level and Curvature*, Econometric Society 2004 North American Winter Meetings, number 432.

[4] Andersen, T.G. and B.E. Sorensen (1996), *GMM Estimation of a Stochastic Volatility Model: A Monte Carlo Study*, Journal of Business and Economic Statistics, vol. 14, pages 328–352.

[5] Arnold, T. and T.F. Crack (1999), *A Practical Guide to GMM*, Working Paper Series.

[6] Bakshi, G., C. Cao and Z. Chen (1997), *Empirical Performance of Alternative Option Pricing Models*, Journal of Finance, vol. 52, pages 2003–2049.

[7] Bates, D.S. (1988), *Pricing options on jump-diffusion processes*, Rodney L. White Center Working paper 37-88, Wharton School.

[8] Bates, D.S. (1996), *Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in the PHLX Deutschemark Options*, Review of Financial Studies, vol. 9, pages 69–107.

[9] Bates, D.S. (2000), *Post-87 Crash Fears in S&P 500 Futures Options Market*, Journal of Econometrics, vol. 94, pages 181–238.

[10] Black F. and M. Scholes (1973), *The Pricing of Options and Corporate Liabilities*, Journal of Political Economy, vol. 81, pages 637–54.

[11] Chernov, M., A.R. Gallant, E. Ghysels and G. Tauchen (2003), *Alternative Models for Stock Price Dynamics*, Journal of Econometrics, vol. 116, pages 225–257.

[12] Cliff, M.T. (2003), *GMM and MINZ program libraries for MATLAB*, Working Paper, Virginia Tech, available for download at `http://sites.google.com/site/mcliffweb/`.

[13] Cont, R. (2001) *Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues*, Quantitative Finance, vol. 1, pages 223–236.

[14] Cox, J.C., J.E. Ingersoll and S. A. Ross (1985), *The Theory of the Term Structure of Interest Rates*, Econometrica, vol. 53, pages 385–407.

[15] Duffie, D., J. Pan and K.J. Singleton (2000), *Transform Analysis and Asset Pricing for Affine Jump-Diffusions*, Econometrica, vol. 68, pages 1343–1376.

[16] Goltz, F., L. Martellini and K. Simsek (2005), *Structured Forms of Investment Strategies in Institutional Investors Portfolios, Benefits of Dynamic Asset Allocation Through Buy-and-Hold Investment in Derivatives*, An Edhec Risk and Asset Management Research Centre Publication.

[17] Hansen, L.P. (1982), *Large Sample Properties of Generalized Method of Moments Estimators*, Econometrica, vol. 50, pages 1029–1054.

[18] Heston, S.I. (1993), *A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options*, Review of Financial Studies, vol. 6, pages 327–344.

[19] Jacquier, E., N.G. Polson and, P.E. Rossi (1994), *Bayesian Analysis of Stochastic Volatility Models*, Journal of Business and Economic Statistics, vol. 12, pages 371–417.

[20] Jiang, G.J. (2002), *Testing Option Pricing Models with Stochastic Volatility, Random Jumps and Stochastic Interest Rates*, International Review of Finance, Vol. 3, pages 233–272.

[21] Kahl, C. and P. Jäckel (2006), *Fast Strong Approximation Monte Carlo Schemes for Stochastic Volatility Models*, Quantitative Finance, vol. 6, pages 513–536.

[22] Liu, J. and J. Pan (2003), *Dynamic Derivative Strategies*, Journal of Financial Economics, vol. 69, pages 401–430.

[23] Newey, W.K. and K.D. West (1987), *A Simple Positive Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix*, Econometrica, vol. 55, pages 703–708.

[24] Pan, J. (2002), *The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study*, Journal of Financial Economics, vol. 63, pages 3–50.

[25] Scott, L. (1997), *Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Application of Fourier Inversion Methods*, Mathematical Finance, vol. 7, pages 413–426.

[26] Westermark. N. (2009), *Barrier Option Pricing*, Degree Project in Mathematics, Royal Institute of Technology (KTH).

Appendix A

Derivation of charachteristic functions

Assume a Data Generating Process that follows the stochastic differential equation (SDE) 2.1. The logarithm of S is then described by the following SDE

$$d\ln S_t = (r_t - d + \lambda\mu + \mu_0)\,dt + \sqrt{V_t}dW_t^S + \ln J_t dq_t(\lambda)$$

We introduce the stochastic process X_t described by $X_t = \ln S_t - \int_0^t \ln J_s\, dq_s$. Note that X_t is independent of $\int_0^t \ln J_s\, dq_s$ and is described by the following SDE.

$$dX_t = (r_t - d + \lambda\mu + \mu_0)\,dt + \sqrt{V_t}dW_t^S$$

Let

$$f^T(x, v, r, t) = \mathrm{E}\left[\exp\left(i\phi_1 X_T + i\phi_2 V_T + i\phi_3 r_T\right)|X_t = x, V_t = v, r_t = r\right]$$

Further, we follow Duffie, Pan and Singleton (2002) and use the following ansatz

$$f^T(x, v, r, t) = \exp\left(C(\phi_1, \phi_2, \phi_3, \Delta) + (i\phi_2 + D(\phi_1, \phi_2, \Delta))\,v + (i\phi_3 B(\phi_1, \phi_3, \Delta))\,r + i\phi_1 x\right)$$

where $\Delta = T - t$.
Now, applying the Itô formula to $f^T(X_t, V_t, r_t, t)$, using the martingale property of conditional expectations and the terminal condition $f^T(x, v, r, T) = \exp\left(i\phi_1 x + i\phi_2 v + i\phi_3 r\right)$ we have the following differential equations

$$\begin{aligned}
\frac{\partial C}{\partial \Delta} &= i\phi_1\left(\mu_0 - \lambda\mu - d\right) + \left(i\phi_2 + D\right)\kappa\gamma + \left(i\phi_3 + B\right)\beta\alpha \\
\frac{\partial D}{\partial \Delta} &= -\frac{1}{2}\sigma_V^2\phi_2^2 - \phi_1\phi_2\sigma_V\rho - \kappa i\phi_2 - \frac{1}{2}\phi_1^2 + \left(\sigma_V^2 i\phi_2 + i\phi_1\sigma_V\rho - \kappa\right)D + \frac{1}{2}\sigma_V^2 D^2 \\
\frac{\partial B}{\partial \Delta} &= i\phi_1 - \frac{1}{2}\sigma_r^2\phi_3^2 - i\phi_3\beta + \left(\sigma_r^2 i\phi_3 - \beta\right)B + \frac{1}{2}\sigma_r^2 B^2
\end{aligned}$$

which are solved by the following functions

$$\begin{aligned}
&C(\phi_1, \phi_2, \phi_3, \Delta) = i\phi_1\Delta\left(\mu_0 - \lambda\mu - d\right) + i\phi_2\Delta\kappa\gamma + \frac{\kappa\gamma}{\sigma_V^2}\left(\kappa - i\phi_1\rho\sigma_V - i\phi_2\sigma_V^2 - h(\phi_1)\right)\Delta \\
&- 2\frac{\kappa\gamma}{\sigma_V^2}\ln\left(\frac{1 - g(\phi_1, \phi_2)\,e^{-h(\phi_1)\Delta}}{1 - g(\phi_1, \phi_2)}\right) + i\phi_3\Delta\beta\alpha + \frac{\beta\alpha}{\sigma_r^2}\left(\beta - i\phi_3\sigma_r^2 - k(\phi_1)\right)\Delta \\
&- 2\frac{\beta\alpha}{\sigma_r^2}\ln\left(\frac{1 - l(\phi_1, \phi_3)\,e^{-k(\phi_1)\Delta}}{1 - l(\phi_1, \phi_3)}\right) \\
&D(\phi_1, \phi_2, \Delta) = \frac{\left(\kappa - i\phi_1\rho\sigma_V - i\phi_2\sigma_V^2 - h(\phi_1)\right)\left(1 - e^{-h(\phi_1)\Delta}\right)}{\sigma_V^2\left(1 - g(\phi_1, \phi_2)\,e^{-h(\phi_1)\Delta}\right)} \\
&B(\phi_1, \phi_3, \Delta) = \frac{\left(\beta - i\phi_3\sigma_r^2 - k(\phi_1)\right)\left(1 - e^{-k(\phi_1)}\right)}{\sigma_r^2\left(1 - l(\phi_1, \phi_3)\,e^{-k(\phi_1)}\right)}
\end{aligned}$$

where

$$
\begin{aligned}
h(\phi_1) &= \sqrt{(\kappa - i\phi_1\rho\sigma_V)^2 + \sigma_V^2\phi_1^2} \\
g(\phi_1,\phi_2) &= \frac{\kappa - i\phi_1\rho\sigma_V - i\phi_2\sigma_V^2 - h(\phi_1)}{\kappa - i\phi_1\rho\sigma_V - i\phi_2\sigma_V^2 + h(\phi_1)} \\
k(\phi_1) &= \sqrt{\beta^2 - 2i\phi_1\sigma_r^2} \\
l(\phi_1,\phi_3) &= \frac{\beta - i\phi_3\sigma_r^2 - k(\phi_1)}{\beta - i\phi_3\sigma_r^2 + k(\phi_1)}
\end{aligned}
$$

Furthermore,

$$
\begin{aligned}
&\mathrm{E}\left[\exp\left(i\phi_1 X_{t+\Delta} + i\phi_2 V_{t+\Delta} + i\phi_3 r_{t+\Delta}\right)\exp\left(-i\phi_1 X_t\right)|X_t, V_t, r_t\right] = \\
&\exp\left(C(\phi_1,\phi_2,\phi_3,\Delta) + (i\phi_2 + D(\phi_1,\phi_2,\Delta))V_t + (i\phi_3 B(\phi_1,\phi_3,\Delta))r_t\right)
\end{aligned}
$$

is obiously independent of X_t.

It is easy to show, using the law of iterated expectations, that

$$
\mathrm{E}\left[\exp\left(i\phi_1\left(\int_0^{t+\Delta}\ln J_s\,dq_s - \int_0^t \ln J_s\,dq_s\right)\right)\right] = \exp\left(\Delta\lambda\left(e^{i\phi_1\mu_J - 1/2\phi_1^2\sigma_J^2} - 1\right)\right)
$$

Hence,

$$
\begin{aligned}
&\psi\left(\Delta\ln S_{t+\Delta}, V_{t+\Delta}, r_{t+\Delta};\phi_1,\phi_2,\phi_3|V_t,r_t\right) = \\
&\mathrm{E}\left[\exp\left(i\phi_1\Delta\ln S_{t+\Delta} + i\phi_2 V_{t+\Delta} + i\phi_3 r_{t+\Delta}\right)|V_t, r_t\right] = \\
&\exp\left(C(\phi_1,\phi_2,\phi_3,\Delta) + (i\phi_2 + D(\phi_1,\phi_2,\Delta))V_t + (i\phi_3 + B(\phi_1,\phi_3,\Delta))r_t\right) \\
&\exp\left(\Delta\lambda\left(e^{i\phi_1\mu_J - 1/2\phi_1^2\sigma_J^2} - 1\right)\right)
\end{aligned}
$$

Now, using that V_t follows a Gamma distribution with density function $f_{V_t}(x) = \frac{\theta^p}{\Gamma(p)}x^{p-1}e^{-\theta x}$, where $\theta = \frac{2\kappa}{\sigma_v^2}$, $p = \frac{2\kappa\gamma}{\sigma_v^2}$ and that r_t follows a Gamma distribution with density function $f_{r_t}(x) = \frac{\theta^p}{\Gamma(p)}x^{p-1}e^{-\theta x}$, where $\theta = \frac{2\beta}{\sigma_r^2}$, $p = \frac{2\beta\alpha}{\sigma_v^2}$ we have,

$$
\begin{aligned}
&\psi\left(\Delta\ln S_{t+\Delta}, r_{t+\Delta};\phi_1,\phi_3\right) = \\
&\mathrm{E}\left[\exp\left(i\phi_1\Delta\ln S_{t+\Delta} + i\phi_3 r_{t+\Delta}\right)\right] = \\
&\mathrm{E}\left[\psi\left(\Delta\ln S_{t+\Delta}, V_{t+\Delta}, r_{t+\Delta};\phi_1,0,\phi_3|V_t,r_t\right)\right] = \\
&\exp\left(C(\phi_1,0,\phi_3,\Delta) - \frac{2\kappa\gamma}{\sigma_V^2}\ln\left(1 - \frac{\sigma_V^2 D(\phi_1,0,\Delta)}{2\kappa}\right) - \frac{2\beta\alpha}{\sigma_r^2}\ln\left(1 - \frac{\sigma_r^2\left(i\phi_3 + B(\phi_1,\phi_3,\Delta)\right)}{2\beta}\right)\right) \\
&\exp\left(\Delta\lambda\left(e^{i\phi_1\mu_J - 1/2\phi_1^2\sigma_J^2} - 1\right)\right)
\end{aligned}
$$

Finally for some $\tau \geq \Delta$,

$$
\begin{aligned}
&\psi\left(\Delta \ln S_{t+\tau+\Delta}, \ln S_{t+\Delta}; \varphi_1, \varphi_2\right) = \mathrm{E}\left[\exp\left(i\varphi_1 \Delta \ln S_{t+\tau+\Delta} + i\varphi_2 \Delta \ln S_{t+\Delta}\right)\right] = \\
&\mathrm{E}\left[\mathrm{E}\left[\exp\left(i\varphi_1 \Delta \ln S_{t+\tau+\Delta} + i\varphi_2 \Delta \ln S_{t+\Delta}\right) | V_{t+\tau}, r_{t+\tau}\right]\right] = \\
&\mathrm{E}\left[\exp\left(i\varphi_2 \Delta \ln S_{t+\Delta}\right) \mathrm{E}\left[\exp\left(C\left(\varphi_1, 0, 0, \Delta\right) + D\left(\varphi_1, 0, \Delta\right) V_{t+\tau} + B\left(\varphi_1, 0, \Delta\right) r_{t+\tau}\right) | V_{t+\Delta}, r_{t+\Delta}\right]\right] \\
&\exp\left(\Delta\lambda\left(e^{i\varphi_1 \mu_J - 1/2\varphi_1^2 \sigma_J^2} - 1\right)\right) = \\
&\exp\left(C\left(\varphi_1, 0, 0, \Delta\right) + C\left(0, -iD\left(\varphi_1, 0, \Delta\right), -iB\left(\varphi_1, 0, \Delta\right), \tau - \Delta\right)\right) \mathrm{E}\left[\exp\left(i\varphi_2 \Delta \ln S_{t+\Delta} + D^* V_{t+\Delta} + B^* r_{t+\Delta}\right)\right] \\
&\exp\left(\Delta\lambda\left(e^{i\varphi_1 \mu_J - 1/2\varphi_1^2 \sigma_J^2} - 1\right)\right) = \\
&\exp\left(C\left(\varphi_1, 0, 0, \Delta\right) + C\left(0, -iD\left(\varphi_1, 0, \Delta\right), -iB\left(\varphi_1, 0, \Delta\right), \tau - \Delta\right) + C\left(\varphi_2, -iD^*, -iB^*, \Delta\right)\right) \\
&\left(1 - \frac{\sigma_V^2\left(D^* + D\left(\varphi_2, -iD^*, \Delta\right)\right)}{2\kappa}\right)^{2\kappa\gamma/\sigma_V^2} \left(1 - \frac{\sigma_r^2\left(B^* + B\left(\varphi_2, -iB^*, \Delta\right)\right)}{2\beta}\right)^{2\beta\alpha/\sigma_r^2} \\
&\exp\left(\Delta\lambda\left(e^{i\varphi_1 \mu_J - 1/2\varphi_1^2 \sigma_J^2} - 1\right) + \Delta\lambda\left(e^{i\varphi_2 \mu_J - 1/2\varphi_2^2 \sigma_J^2} - 1\right)\right)
\end{aligned}
$$

where,

$$
\begin{aligned}
D^* &= D\left(\varphi_1, 0, \Delta\right) + D\left(0, -iD\left(\varphi_1, 0, \Delta\right), \tau - \Delta\right) \\
B^* &= B\left(\varphi_1, 0, \Delta\right) + B\left(0, -iB\left(\varphi_1, 0, \Delta\right), \tau - \Delta\right)
\end{aligned}
$$

Appendix B

Monte Carlo verification of analytical moments

In order to verify the analytical moment conditions as well as to ensure there are no errors in the implementation of the moments or the simulation scheme we simulate 100000 returns using 300 increments per return. Figure B.1 shows that although there is convergence, we require a large amount of simulations. After altering the number of increments as well as the number of simulations we reach the conclusion that the stochastic error in our simulation model is fairly large and that the discretization error is very small in comparison.

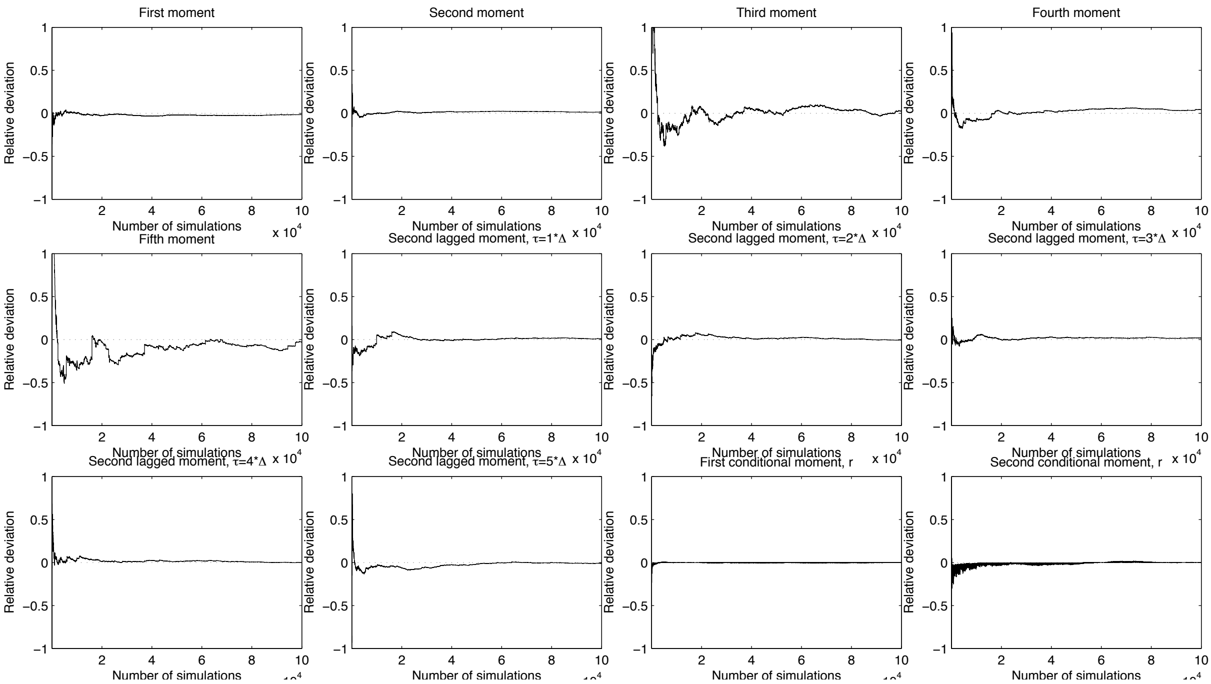

Figure B.1: Relative deviation of Monte Carlo approximation from exact moments.

Appendix C

Option pricing formula

We use the results of Jiang (2002) who derive the following option pricing formula given the risk-neutral dynamics described by Equation 2.2. We will not include the complete derivation, please refer to Jiang (2002) for further details. $C(t, \tau, S_t, K, r_t, V_t)$ denotes the price of a european call option at time t with time to maturity τ and strike price K.

$$C(t, \tau, S_t, K, r_t, V_t) = S_t \Pi_1 (t, \tau, S_t, K, r_t, V_t) - KB(t, \tau) \Pi_2 (t, \tau, S_t, K, r_t, V_t)$$

where

$$\Pi_j (t, \tau, S_t, K, r_t, V_t) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re \left(\frac{e^{-i\phi \ln K} f_j (t, \tau, S_t, K, r_t, V_t, \phi)}{i\phi} \right) d\phi$$

$$f_1 (t, \tau, S_t, K, r_t, V_t, \phi) = \exp \left(-\frac{\gamma_r}{\sigma_r^2} \left(2 \ln \left(1 - \frac{\left(1 - e^{-\xi_r \tau}\right)(\xi_r - \beta^*)}{2\xi_r} \right) (\xi_r - \beta^*) \tau \right) \right)$$

$$\exp \left(-\frac{\gamma_v}{\sigma_v^2} \left(2 \ln \left(1 - \frac{\left(1 - e^{-\xi_v \tau}\right)(\xi_v - \kappa_v + (1 + i\phi) \rho\sigma_v)}{2\xi_v} \right) \right) \right)$$

$$\exp \left(-\frac{\gamma_v}{\sigma_v^2} (\xi_v - \kappa^* + (1 + i\phi) \rho\sigma_v) \tau + i\phi \ln S_t + \frac{2i\phi \left(1 - e^{-\xi_r \tau}\right)}{2\xi_r - (1 - e^{-\xi_r \tau})(\xi_r - \beta^*)} r_t \right)$$

$$\exp \left(\lambda^* \tau (1 + \mu^*) \left((1 + \mu^*)^{i\phi} e^{i\phi(1+i\phi)\sigma_J^2/2} - 1 \right) - i\phi (\lambda^* \mu^* + d) \tau \right)$$

$$\exp \left(\frac{i\phi (i\phi - 1) \left(1 - e^{-\xi_v \tau}\right)}{2\xi_v - (1 - e^{-\xi_v \tau})(\xi_v - \kappa_v + (1 + i\phi) \rho\sigma_v)} V_t \right)$$

$$f_2 (t, \tau, S_t, K, r_t, V_t, \phi) \exp \left(-\frac{\gamma_r}{\sigma_r^2} \left(2 \ln \left(1 - \frac{\left(1 - e^{-\xi_r^* \tau}\right)(\xi_r^* - \beta^*)}{2\xi_r^*} \right) (\xi_r^* - \beta^*) \tau \right) \right)$$

$$\exp \left(-\frac{\gamma_v}{\sigma_v^2} \left(2 \ln \left(1 - \frac{\left(1 - e^{-\xi_v^* \tau}\right)(\xi_v^* - \kappa_v + i\phi\rho\sigma_v)}{2\xi_v^*} \right) \right) \right)$$

$$\exp \left(-\frac{\gamma_v}{\sigma_v^2} (\xi_v^* - \kappa^* + i\phi\rho\sigma_v) \tau + i\phi \ln S_t - \ln B + (t, \tau) \frac{2 (i\phi - 1) \left(1 - e^{-\xi_r^* \tau}\right)}{2\xi_r^* - (1 - e^{-\xi_r^* \tau})(\xi_r^* - \beta^*)} r_t \right)$$

$$\exp \left(\lambda^* \tau (1 + \mu^*) \left((1 + \mu^*)^{i\phi} e^{i\phi(1+i\phi)\sigma_J^2/2} - 1 \right) - i\phi (\lambda^* \mu^* + d) \tau \right)$$

$$\exp \left(\frac{i\phi (i\phi + 1) \left(1 - e^{-\xi_v^* \tau}\right)}{2\xi_v^* - (1 - e^{-\xi_v^* \tau})(\xi_v^* - \kappa_v + i\phi\rho\sigma_v)} V_t \right)$$

where $\xi_r = \sqrt{\beta^{*2} - 2\sigma_r^2 i\phi}$, $\xi_v = \sqrt{(\kappa^* - (1 + i\phi)\rho\sigma_v)^2 - i\phi(1 + i\phi)\sigma_v^2}$, $\xi_r^* = \sqrt{\beta^{*2} - 2\sigma_r^2 (i\phi - 1)}$ and $\xi_v^* = \sqrt{(\kappa^* - i\phi\rho\sigma_v)^2 - i\phi(i\phi - 1)\sigma_v^2}$.