Robustness of Bond Portfolio Optimisation

Divanisha Pillay

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

May 24, 2016

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

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

_____

May 24, 2016

Abstract

Korn and Koziol (2006) apply the Markowitz (1952) mean-variance framework to bond portfolio selection by proposing the use of term structure models to estimate the time-varying moments of bond returns. Duffee (2002) introduces a distinction between completely affine and essentially affine term structure models. A completely affine model uses a market price of risk specification that is proportional to the volatility of the risk factors. However, this assumption of proportionality of the market price of risk contradicts the observed behaviour of bond returns. In response, Duffee (2002) introduces a more flexible essentially affine market price of risk specification by breaking the strict proportionality of the completely affine specification. Essentially affine models better represent the empirical features of bond returns whilst preserving the tractability of completely affine models. However, Duffee and Stanton (2012) find that the increased flexibility of the essentially affine model comes at the expense of real-world parameter estimation. Given these parameter estimation issues, this dissertation investigates whether the difficulty in estimating an essentially affine specification is outweighed by the empirical preferability, and whether, all these issues considered, the Markowitz (1952) approach to bond portfolio optimisation is robust. The results indicate that the superior capability of an essentially affine model to forecast expected returns outweighs real-world parameter estimation issues; and that the estimation and mean-variance optimisation procedures are worthwhile.

Acknowledgements

I would to thank my supervisors, Alex Backwell and doctor Peter Ouwehand, for their guidance, patience and keen interest. As always, I am grateful to my family for their continued support and encouragement.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The Markowitz (1952) mean-variance approach to portfolio optimisation has been widely applied to stock selection. However, little attention has been given to selecting bond portfolios using this framework. Korn and Koziol (2006) suggest that this is largely because, historically, interest rate volatility has been low and, thus, portfolio optimisation seemed unnecessary. However, over the last few decades, changes in interest rates have posed significant risks to bond investments, even ignoring default risk. There is scope for risk diversification using bonds because of the availability of a large number of different bonds with different maturities. In particular, this paper investigates the application of the mean-variance framework in selecting amongst different bond maturities.

The main issue in applying mean-variance optimisation to bond portfolio selection is that the moments of bond returns vary over time. Therefore, naive historical estimation of these moments —namely, the expected return and covariance matrix of returns —assuming stationarity is inadequate. Korn and Koziol (2006) propose using term structure models to estimate these moments. Term structure models can be classified as either equilibrium or no-arbitrage models. The Vasicek (1977) and CIR [1] models are examples of equilibrium models. The success of term structure models in pricing interest rate dependent securities is widely documented in the literature. Term structure models are useful because they capture the variation of bond returns as the state of the world/the state variables as well as time to maturity varies. The use of term structure models in bond portfolio optimisation was originally suggested by Brennan and Schwartz (1980). Wilhelm (1992) specifies a CIR model for bond returns and derives the requisite means, variances and covariances of returns over discrete holding periods. Korn and Koziol (2006) are the first to test such optimised portfolio strategies in an empirical setting. More specifically, they base their study on the German government bond market over 1974 to 2004. Puhle (2008) extends the mean-variance bond portfolio optimisation framework of Wil-

---

[1] Proposed by Cox *et al.* (1985)

helm (1992) to include the Vasicek (1977) and Hull and White (1994) term structure models. Caldeira *et al.* (2012) make use of dynamic heteroskedastic factor models to better capture the persistence in the volatility of bond returns.

Duffie and Kan (1996) characterise the affine class of term structure models by a set of conditions on the risk factors, the short rate and the bond price. In the affine class, bond prices have an exponential-affine form, which implies that the drift, squared diffusion and short rate are affine functions of the risk factors. Duffee (2002) notes that despite the existence of a range of non-affine models, the tractability and richness of the affine class have made these models the focus of much attention in the literature.

Duffee (2002) highlights that when the market price of risk is proportional to the volatility of the risk factors, then there is linearity under both the risk-neutral and real-world measures. In this case, the model is known as a completely affine term structure model. However, the assumption of proportionality of the market price of risk contradicts the observed behaviour of bond returns. The empirical limitations of completely affine term structure models include, firstly, poor forecasting of excess returns on long-dated bonds; and, secondly, the inability to capture the time-varying volatility of interest rates. In response to these empirical limitations, Duffee (2002) and Duarte (2004) introduce essentially affine and semi-affine term structure models, which allow for more flexible price of risk specifications. These models better represent the empirical features of bond returns whilst preserving the tractability of completely affine term structure models.

However, the increased flexibility of the essentially affine and semi-affine models comes at the expense of parameter estimation. Duffee and Stanton (2012) investigate the finite-sample properties of three estimation techniques through Monte Carlo simulations, namely, maximum likelihood, efficient method of moments and the Kalman filter. They find significant biases in the real-world parameter estimates across all estimation techniques.

Given these issues relating to real-world parameter estimation, this dissertation will investigate whether the difficulty in estimating an essentially affine specification is outweighed by the empirical preferability, and whether, all these issues considered, the Markowitz approach to bond portfolio optimisation is robust in terms of performance and estimation of expected returns versus realised returns. The use of an essentially affine price of risk specification results in affine real-world dynamics, whereas a semi-affine price of risk specification does not. This is important for the calculation of bond moments under the real-world measure, which are required as inputs to mean-variance optimisation.

Chapter 2 details affine term structure models along with the completely and

essentially affine price of risk specifications. The Vasicek (1977) model is introduced and the moments of bond returns under the real-world measure are derived. Chapter 3 outlines the issues related to parameter estimation when a more flexible price of risk is used and provides a description of the Kalman filter. In addition, the accuracy of the Kalman filter in parameter identification is investigated through a Monte Carlo simulation experiment. Chapter 4 details the bond portfolio optimisation procedure and provides results from the testing of the optimised portfolios. Finally, Chapter 5 concludes.

Chapter 2

Affine Term Structure Models: Market Price of Risk Specification

2.1 Completely Affine Versus Essentially Affine

This section summarises the affine term structure model framework developed by Duffie and Kan (1996). Assume there are n factors, $X(t) \equiv (X_1(t), \ldots, X_n(t))'$, and n Brownian motions, $\tilde{W}(t) \equiv (\tilde{W}_1(t), \ldots, \tilde{W}_n(t))'$. The dynamics of the factors under the equivalent martingale measure, $\mathbb{Q}$, are given by

$$dX(t) = \kappa(\theta - X(t))dt + \Sigma S(t)d\tilde{W}(t),$$

where κ and Σ are $n \times n$ matrices and θ is an $n \times 1$ vector. $S(t)$ is a diagonal matrix with typical element

$$S_{(ii)}(t) \equiv \sqrt{\alpha_i + \beta_i' X(t)},$$

where β_i is an $n \times 1$ vector and α_i a scalar. Parameter restrictions required for $\alpha_i + \beta_i' X(t)$ to be nonnegative for all i and all feasible $X(t)$ can be found in Dai and Singleton (2000). The instantaneous spot rate (short rate) is an affine function of these factors:

$$r(t) = \delta_0 + \delta' X(t),$$

where δ_0 is a scalar and δ is a $n \times 1$ vector.

Let $P(X(t), \tau)$ be the time-t price of a zero-coupon bond maturing at time $t + \tau$. Duffie and Kan (1996) prove that the bond price has an exponential affine form

$$P(X(t), \tau) = e^{A(\tau) - B(\tau)' X(t)},$$

where $A(\tau)$ is a scalar and $B(\tau)$ is an $n \times 1$ vector. Therefore, the bond's yield is an affine function of the factors:

$$Y(X(t), \tau) = \frac{1}{\tau}[-A(\tau) + B(\tau)' X(t)].$$

The functions $A(\tau)$ and $B(\tau)$ are the solutions to a system of ordinary differential equations (ODEs) provided in Dai and Singleton (2000):

$$\frac{dA(\tau)}{d\tau} = -\theta'\kappa'B(\tau) + \frac{1}{2}\sum_{i=1}^{n}[\Sigma'B(\tau)]_i^2\alpha_i - \delta_0, \qquad (2.1)$$

$$\frac{dB(\tau)}{d(\tau)} = -\kappa'B(\tau) - \frac{1}{2}\sum_{i=1}^{n}[\Sigma'B(\tau)]_i^2\beta_i + \delta, \qquad (2.2)$$

with initial conditions $A(0) = 0$ and $B(0) = 0_{n\times 1}$.

In order to move from $\mathbb{Q}$ to the real-world measure $\mathbb{P}$, Duffee (2002) specifies the dynamics of the state price deflator as

$$\frac{d\pi(t)}{\pi(t)} = -r(t)dt - \Lambda(t)'dW(t),$$

where $W(t)$ is an $n \times 1$ $\mathbb{P}$-Brownian motion and $\Lambda(t)$ is an $n \times 1$ vector representing the market price of risk. By Girsanov's theorem,

$$\tilde{W}(t) = W(t) + \int_0^t \Lambda(s)ds$$

is a $\mathbb{Q}$-Brownian motion. The $\mathbb{P}$-dynamics of $X(t)$ are, therefore, given by

$$dX(t) = \kappa(\theta - X(t))dt + \Sigma S(t)\Lambda(t)dt + \Sigma S(t)dW(t).$$

The instantaneous bond-price dynamics are given by

$$\frac{dP(X(t),\tau)}{P(X(t),\tau)} = (r(t) + e_\tau(t))dt + v_\tau(t)dW(t),$$

where $e_\tau(t)$ is the instantaneous expected excess bond return and $v_\tau(t)$ is the instantaneous bond return volatility. Duffee (2002) provides expressions for $e_\tau(t)$ and $v_\tau(t)$:

$$e_\tau(t) = -B(\tau)'\Sigma S(t)\Lambda(t),$$
$$v_\tau(t) = -B(\tau)'\Sigma S(t).$$

Therefore, changes in expected excess bond returns over time are caused by changes in both the volatility matrix $S(t)$ and the market price of risk vector $\Lambda(t)$. There is a strong connection between excess bond returns and bond volatility due to their mutual dependence on factor volatility.

The *completely affine* price of risk specification of Fisher and Gilles (1996) and Dai and Singleton (2000) is

$$\Lambda(t) = S(t)\lambda_1,$$

where λ_1 is an $n \times 1$ vector.

Duffee (2002) characterises the *essentially affine* class of term structure models by first defining $S^-(t)$, a diagonal matrix with elements

$$S_{(ii)}^-(t) = \begin{cases} (\alpha_i + \beta_i' X(t))^{-\frac{1}{2}}, & \text{if } \inf(\alpha_i + \beta_i' X(t)) > 0; \\ 0, & \text{otherwise.} \end{cases}$$

This definition of $S^-(t)$ ensures that its elements do not tend to infinity as the associated elements of $S(t)$ tend to zero.

The form of the essentially affine price of risk is

$$\Lambda(t) = S(t)\lambda_1 + S^-(t)\lambda_2 X(t),$$

where λ_2 is an $n \times n$ matrix.

For both the completely affine and essentially affine cases, $S(t)\Lambda(t)$ is affine in $X(t)$. Therefore, $X(t)$ has affine dynamics under both $\mathbb{Q}$ and $\mathbb{P}$. Affine $\mathbb{P}$-dynamics are desirable because they allow closed-form solutions for the conditional mean and variance of discretely sampled bond yields. "Completely affine" refers to $\Lambda(t)'\Lambda(t)$, the instantaneous variance of the log state price deflator, also being affine in $X(t)$. In the essentially affine case, $\Lambda(t)'\Lambda(t)$ is not affine in $X(t)$ when $\lambda_2 \neq 0$. This is a less important property since the variance of the state price deflator does not influence bond prices.

Duarte (2004) highlights the following empirical limitations of completely affine term structure models: (i) poor forecasting of excess returns on long-dated treasury bonds; and (ii) the inability to capture the time-varying volatility of interest rates. Duffee (2002) finds that the essentially affine class addresses these failures, whilst maintaining tractability, by removing the strict proportionality between the price of interest rate risk and interest rate volatility. This allows for independent movements of the price of risk and results in improved forecasting of future yields. In addition, the essentially affine specification allows $\Lambda(t)$ to change sign over time with the shape of the term structure.

2.2 Vasicek Model Specification

The Vasicek model falls within the affine class of term structure models. Vasicek (1977) assumed that the short rate follows an Ornstein-Uhlenbeck process with constant coefficients. Assume a Vasicek-type multi-factor model where the factors $X_i(t)$, $i = 1, \ldots, n$, are modelled as mutually independent Ornstein-Uhlenbeck processes, with risk-neutral or equivalent-martingale ($\mathbb{Q}$) dynamics given by

$$dX_i(t) = \kappa_i(\theta_i - X_i(t))dt + \sigma_i d\tilde{W}_i(t).$$

The speed of mean reversion of factor i is given by κ_i, θ_i is the long-term mean of factor i, and σ_i represents the local volatility of the factor i increment. Assume the short rate equals the sum of the factors:

$$r(t) = \sum_{i=1}^{n} X_i(t).$$

Brigo and Mercurio (2007) integrate the stochastic differential equation (SDE) to obtain, for $s \le t$,

$$X_i(t) = X_i(s)e^{-\kappa_i(t-s)} + \theta_i \left(1 - e^{-\kappa_i(t-s)}\right) + \sigma_i \int_s^t e^{-\kappa_i(t-u)} d\tilde{W}_i(u),$$

so that $X_i(t)$ conditional on $\mathcal{F}_s$ is normally distributed under $\mathbb{Q}$ with mean and variance:

$$E^{\mathbb{Q}}\left[X_i(t)|\mathcal{F}_s\right] = X_i(s)e^{-\kappa_i(t-s)} + \theta_i \left(1 - e^{-\kappa_i(t-s)}\right),$$

$$Var^{\mathbb{Q}}\left[X_i(t)|\mathcal{F}_s\right] = \frac{\sigma_i^2}{2\kappa_i} \left(1 - e^{-2\kappa_i(t-s)}\right).$$

Therefore, $r(t)$ is also normally distributed with mean and variance:

$$E^{\mathbb{Q}}\left[r(t)|\mathcal{F}_s\right] = \sum_{i=1}^{n} \left[X_i(s)e^{-\kappa_i(t-s)} + \theta_i \left(1 - e^{-\kappa_i(t-s)}\right)\right],$$

$$Var^{\mathbb{Q}}\left[r(t)|\mathcal{F}_s\right] = \sum_{i=1}^{n} \frac{\sigma_i^2}{2\kappa_i} \left(1 - e^{-2\kappa_i(t-s)}\right)$$

Bond prices in the Vasicek model have an exponential affine form, $P(X(t), \tau) = e^{A(\tau) - B(\tau)'X(t)}$. One can verify that Equations (2.1) and (2.2) with the appropriate parameters, $\alpha_i = 1$, β_i an $n \times 1$ zero vector, $\delta_0 = 0$ and δ an $n \times 1$ vector of ones, are satisfied by

$$A(\tau) = \sum_{i=1}^{n} \left[(\theta_i - \frac{\sigma_i^2}{2\kappa_i^2})(B_i(\tau) - \tau) - \frac{\sigma_i^2 B_i^2(\tau)}{4\kappa_i}\right],$$

$$B_i(\tau) = \frac{1}{\kappa_i}(1 - e^{-\kappa_i \tau}).$$

2.3 Market Price of Risk Specification

The completely affine price of risk specification investigated is

$$\Lambda_i(t) = \frac{\lambda_i}{\sigma_i}.$$

Therefore, in the completely affine case, the dynamics of the factors under the real-world measure, $\mathbb{P}$, are given by

$$dX_i(t) = \kappa_i \left(\theta_i + \frac{\lambda_i}{\kappa_i} - X_i(t) \right) dt + \sigma_i dW_i(t).$$

Under $\mathbb{P}$, $X_i(t)$ conditional on $\mathcal{F}_s$ is normally distributed with mean and variance:

$$E^{\mathbb{P}}\left[X_i(t)|\mathcal{F}_s\right] = X_i(s)e^{-\kappa_i(t-s)} + \left(\theta_i + \frac{\lambda_i}{\kappa_i} \right) \left(1 - e^{-\kappa_i(t-s)} \right),$$

$$Var^{\mathbb{P}}\left[X_i(t)|\mathcal{F}_s\right] = \frac{\sigma_i^2}{2\kappa_i} \left(1 - e^{-2\kappa_i(t-s)} \right).$$

The essentially affine price of risk specification takes the form investigated by Duffee and Stanton (2012):

$$\Lambda_i(t) = \frac{\lambda_{i1} + \lambda_{i2}X_i(t)}{\sigma_i}.$$

In this case, the $\mathbb{P}$-dynamics of the factors are given by

$$dX_i(t) = (\kappa_i\theta_i + \lambda_{i1} - (\kappa_i - \lambda_{i2})X_i(t))dt + \sigma_i dW_i(t).$$

With this specification, $X_i(t)$ remains conditionally normally distributed under $\mathbb{P}$ with mean and variance:

$$E^{\mathbb{P}}\left[X_i(t)|\mathcal{F}_s\right] = X_i(s)e^{-(\kappa_i-\lambda_{i2})(t-s)} + \left(\frac{\kappa_i\theta_i + \lambda_{i1}}{\kappa_i - \lambda_{i2}} \right) \left(1 - e^{-(\kappa_i-\lambda_{i2})(t-s)} \right),$$

$$Var^{\mathbb{P}}\left[X_i(t)|\mathcal{F}_s\right] = \frac{\sigma_i^2}{2(\kappa_i - \lambda_{i2})} \left(1 - e^{-2(\kappa_i-\lambda_{i2})(t-s)} \right).$$

2.4 Moments of Bond Returns

For the investor's portfolio optimisation problem consisting of M bonds, the expected returns μ_k, $k = 1, \ldots, M$, of every risky bond and covariance of returns $\Omega = \{s_{k,l}^2\}$, $k, l = 1, \ldots, M$, are required under the real-world measure. Let $P(t, T_k)$ [1] be the price of a bond at time t that matures at time T_k. The expected return of bond k, μ_k, and covariance of bond k and bond l's returns, $s_{k,l}^2$, over the period $t = 0$ to T are calculated as

$$\mu_k = \frac{E_0^{\mathbb{P}}\left[P(T, T_k)\right]}{P(0, T_k)} - 1,$$

$$s_{k,l}^2 = Cov_0^{\mathbb{P}}\left[\frac{P(T, T_k)}{P(0, T_k)}, \frac{P(T, T_l)}{P(0, T_l)} \right].$$

---

[1] Note the change in notation from earlier sections

Korn and Koziol (2006) allow for real-world bond pricing errors at the planning horizon:

$$P(T, T_k) = e^{A(T_k - T) - \sum_{i=1}^{n} X_i(T) B_i(T_k - T) + \epsilon_{T_k}(T)},$$

where $\epsilon_{T_k}(T)$ is the error term associated with maturity T_k. $\epsilon_{T_k}(T)$ is assumed to be normally distributed with zero mean and variance $s^2(\epsilon_{T_k})$. These pricing errors account for the less than perfect model specification and are handled naturally by the Kalman filter in the next section. Maximum likelihood estimation can be used to estimate the pricing error variances along with the other parameters. Error terms referring to different maturities are assumed to be mutually independent. Allowance is made for pricing errors in this dissertation as well.

Since bond prices are log-normally distributed in the Vasicek model, closed-form solutions exist for the expected return μ_k, variance $s^2_{k,k}$ and covariance $s^2_{k,l}$, which are provided in Korn and Koziol (2006):

$$\mu_k = \frac{e^{M^{(1)}(T_k) + \frac{1}{2} S^{(1)}(T_k)^2}}{P(0, T_k)} - 1,$$

$$s^2_{k,k} = \frac{e^{2M^{(1)}(T_k) + S^{(1)}(T_k)^2} \left(e^{S^{(1)}(T_k)^2} - 1 \right)}{P(0, T_k)^2},$$

$$s^2_{k,l} = \frac{e^{M^{(2)}(T_k, T_l) + \frac{1}{2} S^{(2)}(T_k, T_l)^2} - e^{M^{(1)}(T_k) + M^{(1)}(T_l) + \frac{1}{2}\left(S^{(1)}(T_k)^2 + S^{(1)}(T_l)^2 \right)}}{P(0, T_k) P(0, T_l)}, \text{ for } k \neq l,$$

where

$$M^{(1)}(T_k) = A(T_k - T) - \sum_{i=1}^{n} \left(E_0^{\mathbb{P}} [X_i(T)] B_i(T_k - T) \right),$$

$$S^{(1)}(T_k) = \sqrt{\sum_{i=1}^{n} \left(Var_0^{\mathbb{P}} [X_i(T)] B_i(T_k - T)^2 \right) + s^2(\epsilon_{T_k})},$$

$$M^{(2)}(T_k, T_l) = A(T_k - T) + A(T_l - T) - \sum_{i=1}^{n} \left(E_0^{\mathbb{P}} [X_i(T)] \left(B_i(T_k - T) + B_i(T_l - T) \right) \right),$$

$$S^{(2)}(T_k, T_l) = \sqrt{\sum_{i=1}^{n} \left(Var_0^{\mathbb{P}} [X_i(T)] \left(B_i(T_k - T) + B_i(T_l - T) \right)^2 \right) + s^2(\epsilon_{T_k}) + s^2(\epsilon_{T_l})}.$$

Chapter 3

Parameter Estimation

3.1 Issues Related to a More Flexible Price of Risk

This chapter focuses on the findings of Duffee and Stanton (2012) regarding the one-factor model case; they observe similar results in a two-factor setting. Using a completely affine price of risk specification, the drift of $X(t)$ under the real-world measure shares common parameters with the drift of $X(t)$ under the risk-neutral measure, namely the speed of mean reversion κ. The drift is determined precisely by both cross-sectional and longitudinal information. However, using an essentially affine specification, the real-world and risk-neutral drifts share no common parameters. Therefore, the parameters $(\kappa\theta + \lambda_1)$ and $(\kappa - \lambda_2)$ are estimated using only the time-series characteristics of bond yields. Duffee and Stanton (2012) find significant finite-sample biases and high variability in the estimates of the parameters determined only under the real-world measure. In contrast, the estimates of the parameters determined under the risk-neutral measure show little to no bias and low variability. After 500 Monte Carlo simulations of maximum likelihood estimation by means of the Kalman filter, Duffee and Stanton (2012) find biases of approximately one standard deviation in the estimate of λ_1 and minus one standard deviation in the estimate of λ_2.

Duffee and Stanton (2012) consider the effects of the bias in parameter estimation for the essentially affine model on expected excess bond returns. A bond with time to maturity τ has a time-t instantaneous expected excess return given by

$$e_\tau(t) = -\frac{1 - e^{-\kappa\tau}}{\kappa} \left(\lambda_1 + \lambda_2 r(t)\right).$$

This expression is readily interpreted. The fraction gives the sensitivity of a bond's log price to instantaneous interest rates. For values of κ near zero, it provides an approximation of the time to maturity of the bond. The expression in parentheses represents the difference between the real-world drift of $r(t)$ and the risk-neutral drift. The higher the real-world drift compared to the risk-neutral drift, the lower

the expected excess bond return. This implies that investors are pricing bonds as though interest rates will increase more steeply (or decrease more quickly) than what is expected under the real-world measure. The failure of the expectations hypothesis to explain the behaviour of Treasury yields is widely documented. For instance, Campbell and Shiller (1991) state that expected excess bond returns are high when the term structure slope is steep. Duffee and Stanton (2012) note that the parameter estimation bias they observe works in favour of the expectations hypothesis. Consider the effect of the negative bias in λ_2. For low values of $r(t)$ (when the slope is steep), the expected excess returns are lower than those implied by the true parameter values.

3.2 The Kalman Filter

This section is based on the description of the Kalman filter contained in Bolder (2001). Duan and Simonato (1999), Lund (1997), Geyer and Pichler (1999), De Jong (1999) and Babbs and Nowman (1999) laid the foundations for the application of the Kalman filter technique within the affine term-structure framework. The Kalman filter is useful in the context of parameter estimation where the underlying state variables or factors are unobservable. The filter uses the information contained in bond prices to identify the underlying state variables. The Kalman filter requires the interest rate model to be cast in state-space form, which comprises the measurement system and the transition system. The measurement system specifies the linear relationship between market zero-coupon rates (obtained from the log of the bond price) and the underlying state variables. The transition system is an unobserved system of equations which expresses the dynamics of the state variables specified under the model. The state-space formulation allows the Kalman filter to make recursive inferences regarding the unobserved state variable values (transition system) based on the observed market zero-coupon rates (measurement system). Finally, the optimal parameter estimates are obtained by maximum likelihood estimation given the recursive inferences.

Assume an n-factor model. Bolder (2001) states that normally only one market zero-coupon rate is needed for each factor in the estimation procedure. However, increasing the number of market zero-coupon rates adds to the cross-sectional information, which improves estimation precision. Assume M zero-coupon rates are used. These rates together with the measurement equation, specified below, form the measurement system. The measurement equation expresses the relationship

between the zero-coupon yield and zero-coupon price:

$$z(t, T_k) = -\frac{\ln P(t, T_k)}{T_k - t} = \frac{-A(T_k - t) + \sum_{i=1}^{n} X_i(t) B_i(T_k - t)}{T_k - t},$$

where $z(t, T_k)$ denotes the zero-coupon yield over the period t to T_k.

Bolder (2001) creates N intervals of equal length over the range $[t, T_k]$. Let $t_j = j \frac{(T_k - t)}{N}$ for $j = 1, \ldots, N$. $\Delta t = t_j - t_{j-1}$ represents the size of the time step. ϵ_{t_j} is the measurement error introduced into the system to account for pricing errors. The measurement equation is given in vector form as

$$z_{t_j} = D + H X_{t_j} + \epsilon_{t_j},$$

where ϵ_{t_j} is the measurement error introduced into the system to account for pricing errors, assumed to be normally distributed with zero mean and covariance matrix R. The vector D and matrices H and R are given by

$$D = \left[\frac{-A(t_{z_i} - t_j)}{t_{z_i} - t_j} \right]_{i=1,\ldots,M},$$

$$H = \begin{bmatrix} \frac{B_1(t_{z_1} - t_j)}{t_{z_1} - t_j} & \cdots & \frac{B_n(t_{z_1} - t_j)}{t_{z_1} - t_j} \\ & \vdots & \\ \frac{B_1(t_{z_M} - t_j)}{t_{z_M} - t_j} & \cdots & \frac{B_n(t_{z_M} - t_j)}{t_{z_M} - t_j} \end{bmatrix},$$

$$R = diag \left[r_i^2 \right]_{i=1,\ldots,M}.$$

Note that t_{z_i} is the time point corresponding to the i^{th} zero-coupon rate. The transition equation describes the discrete-time evolution of $X(t)$ under the real-world measure as a linear function of $X(t)$. The Vasicek SDE can be solved explicitly for $X(t)$ and this solution discretised. In the completely affine case, the transition equation is specified as

$$X_{t_j} = C + F X_{t_{j-1}} + \nu_{t_j},$$

where ν_{t_j} is the innovations process assumed to be normally distributed with zero mean and covariance matrix Q. The vector C and matrices F and Q are given by

$$C = \left[(\theta_i + \frac{\lambda_i}{\kappa_i})(1 - e^{-\kappa_i \Delta t}) \right]_{i=1,\ldots,n},$$

$$F = diag \left[e^{-\kappa_i \Delta t} \right]_{i=1,\ldots,n},$$

$$Q = diag \left[\frac{\sigma_i^2}{2\kappa_i}(1 - e^{-2\kappa_i \Delta t}) \right]_{i=1,\ldots,n}.$$

In the essentially affine case, the transition equation has vector C and matrices F and Q given by

$$C = \left[\frac{\kappa_i \theta_i + \lambda_{i1}}{\kappa_i - \lambda_{i2}} \left(1 - e^{-(\kappa_i - \lambda_{i2})\Delta t}\right) \right]_{i=1,\ldots,n},$$

$$F = diag \left[e^{-(\kappa_i - \lambda_{i2})\Delta t} \right]_{i=1,\ldots,n},$$

$$Q = diag \left[\frac{\sigma_i^2}{2(\kappa_i - \lambda_{i2})} \left(1 - e^{-2(\kappa_i - \lambda_{i2})\Delta t}\right) \right]_{i=1,\ldots,n}.$$

3.2.1 Filtering Steps

This section provides an outline of the steps implemented in the filtering procedure based on the description found in Bolder (2001).

1. State vector initialisation

The unconditional mean and variance of the transition system are used as starting estimates for the state vector. In the completely affine case, the starting values are

$$\mathrm{E}[X_1] = \mathrm{E}[X_1|\mathcal{F}_0] = \left[\theta_i + \frac{\lambda_i}{\kappa_i} \right]_{i=1,\ldots,n},$$

$$\mathrm{Var}[X_1] = \mathrm{Var}[X_1|\mathcal{F}_0] = diag \left[\frac{\sigma_i^2}{2\kappa_i} \right]_{i=1,\ldots,n}.$$

In the essentially affine case, the starting values are

$$\mathrm{E}[X_1] = \mathrm{E}[X_1|\mathcal{F}_0] = \left[\frac{\kappa_i \theta_i + \lambda_{i1}}{\kappa_i - \lambda_{i2}} \right]_{i=1,\ldots,n},$$

$$\mathrm{Var}[X_1] = \mathrm{Var}[X_1|\mathcal{F}_0] = \left[\frac{\sigma_i^2}{2(\kappa_i - \lambda_{i2})} \right]_{i=1,\ldots,n}.$$

2. Observation forecast

The conditional forecast of the measurement equation and its conditional variance are

$$\mathrm{E}[z_{t_j}|\mathcal{F}_{t_{j-1}}] = D + H\mathrm{E}[X_{t_j}|\mathcal{F}_{t_{j-1}}],$$
$$\mathrm{Var}[z_{t_j}|\mathcal{F}_{t_{j-1}}] = H\mathrm{Var}[X_{t_j}|\mathcal{F}_{t_{j-1}}]H' + R.$$

3. Inference update

Given the realised value of the measurement system, z_{t_j}, the prediction error is given by

$$\zeta_{t_j} = z_{t_j} - \mathrm{E}[z_{t_j}|\mathcal{F}_{t_{j-1}}].$$

This prediction error is used to calculate an updated inference about the state vector and its conditional variance as follows:

$$E[X_{t_j}|\mathcal{F}_{t_j}] = E[X_{t_j}|\mathcal{F}_{t_{j-1}}] + K_{t_j}\zeta_{t_j},$$
$$\text{Var}[X_{t_j}|\mathcal{F}_{t_j}] = (I - K_{t_j}H)\text{Var}[X_{t_j}|\mathcal{F}_{t_{j-1}}].$$

where

$$K_{t_j} = \text{Var}[X_{t_j}|\mathcal{F}_{t_{j-1}}]H'\text{Var}[z_{t_j}|\mathcal{F}_{t_{j-1}}]^{-1}.$$

K_{t_j} is referred to as the Kalman gain matrix.

4. State vector forecast

This step involves forecasting the state vector for the subsequent time period based on the previous updated state vector:

$$E[X_{t_{j+1}}|\mathcal{F}_{t_j}] = C + FE[X_{t_j}|\mathcal{F}_{t_j}],$$
$$\text{Var}[X_{t_{j+1}}|\mathcal{F}_{t_j}] = \text{Var}[X_{t_j}|\mathcal{F}_{t_{j-1}}] - F\text{Var}[X_{t_j}|\mathcal{F}_{t_j}]F' + Q.$$

5. Maximum likelihood estimation

The above steps are repeated at each time step in order to generate a time series of underlying state variables. The log-likelihood function can be constructed based on the assumption that the prediction errors are Gaussian:

$$\ell(\rho) = \sum_{j=1}^{N}\ln[(2\pi)^{-\frac{M}{2}}\det(\text{Var}[z_{t_j}|\mathcal{F}_{t_{j-1}}])^{-\frac{1}{2}}e^{-\frac{1}{2}\zeta'_{t_j}\text{Var}[z_{t_j}|\mathcal{F}_{t_{j-1}}]^{-1}\zeta_{t_j}}],$$
$$= -\frac{MN\ln(2\pi)}{2} - \frac{1}{2}\sum_{j=1}^{N}[\ln(\det(\text{Var}[z_{t_j}|\mathcal{F}_{t_{j-1}}])) + \zeta'_{t_j}\text{Var}[z_{t_j}|\mathcal{F}_{t_{j-1}}]^{-1}\zeta_{t_j}].$$

The optimal parameter set for the Vasicek model is that which maximises the log-likelihood. Optimisation involves the use of non-linear numerical techniques.

3.3   Monte Carlo Simulations

The effectiveness of the Kalman filter in parameter identification is tested through a Monte Carlo simulation experiment. Data are simulated based on a known parameter set and the Kalman filter is then used to estimate these parameters. After 50 simulations, the mean and standard deviation of the parameter estimates is calculated in order to determine estimation precision. The results of the estimation for

the one- and two-factor models are displayed in Tables 3.1 and 3.2.

Tab. 3.1: Simulation results for completely affine (CA) and essentially affine (EA) one-factor models

	True Values	CA Estimates	CA Std Dev	EA Estimates	EA Std Dev
θ	0.1	0.10001	0.00012795	0.098859	0.00022378
κ	0.7	0.7	0.00043892	0.70088	0.00032141
σ	0.06	0.060026	0.0010812	0.049675	0.0022361
λ_1	-0.17	-0.17	0.0048249	-0.28862	0.034017
λ_2	-0.1	-	-	-0.19315	0.12355
σ_k	0.001	0.0009989	7.887e-06	0.0010014	6.4787e-06

Table 3.1 shows that the parameter estimates for the completely affine one-factor model are very close to unbiased and have a low variability. In the essentially affine case, the parameters which are identified under the risk-neutral measure, θ,κ and σ are estimated with very little bias and σ_k little bias, despite its error being an order of magnitude larger than that of the former parameters. However, the estimates of the risk premia, λ_1 and λ_2, are strongly negatively biased and show an increased variability.

Tab. 3.2: Simulation results for completely and essentially affine two-factor models

	True Values	CA Estimates	CA Std Dev	EA Estimates	EA Std Dev
θ_1	0.1	0.085203	0.0043991	0.084452	0.029093
θ_2	0.07	0.084944	0.0043865	0.075749	0.029042
κ_1	0.7	0.70342	0.010056	0.7124	0.0099588
κ_2	0.4	0.40083	0.0024875	0.39922	0.0030625
σ_1	0.06	0.059925	0.001919	0.023693	0.001795
σ_2	0.05	0.050742	0.0012833	0.018423	0.001344
λ_{11}	-0.17	-0.16823	0.007443	-0.18091	0.022844
λ_{12}	-0.1	-	-	-0.20018	0.046381
λ_{21}	-0.2	-0.20216	0.0073874	-0.13472	0.18708
λ_{22}	-0.15	-	-	-0.17057	0.15455
σ_k	0.001	0.0010001	6.3767e-06	0.00099671	1.2175e-05

The parameter estimates for a two-factor model are less precise than that of a one-factor model. In particular, there is difficulty in distinguishing between θ_1 and θ_2 in both the completely and essentially affine cases. For the completely affine model, the rest of the parameter estimates are close to unbiased and show little variability. However, the estimates for the essentially affine volatility and risk pre-

mia parameters are negatively biased. The estimates for λ_{i2}, $i = 1, \ldots, n$, are particularly variable as documented in Duffee and Stanton (2012). Also documented in Duffee (2002) is the trade-off between forecasting future yields and estimating interest rate volatility.

Figure 3.1 shows the effectiveness of the Kalman filter in state identification. Specifically, it shows filtered versus simulated underlying state variables for an essentially affine three-factor model. Observe that, in some instances, the Kalman filter has difficulty in distinguishing between the state variables corresponding to each factor. However, the sum of the filtered state variables corresponding to each factor results in a path which lies very close to the sum of the simulated state variables. Therefore, equating to the sum of factors means that the yield curve is being matched almost exactly.

Fig. 3.1: Filtered state variables for an essentially affine three-factor model
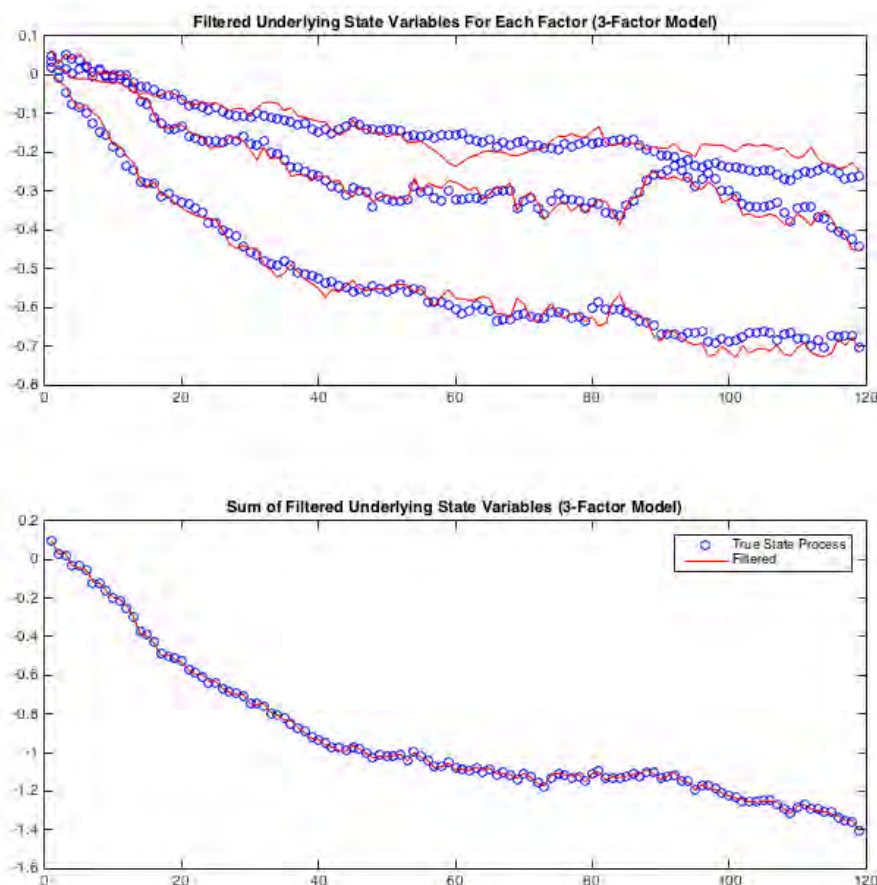
Chapter 4

Bond Portfolio Optimisation

4.1   Overview of Estimation Procedure

The data consists of daily continuously compounded South African yield curves bootstrapped from swap instruments over the period 2 January 2004 to 26 March 2015. Swap market data is chosen over bond market data because: (i) the number of maturities/'bonds' (this dissertation fixes a number of maturities for the analysis, rather than following particular bonds as they approach maturity); and (ii) the South African government bond market is quite illiquid, while the swap market is deep and liquid. An investment period of one year is chosen. A four-year period immediately preceding the investment period is used to estimate model parameters. On each day in the estimation period, 16 different zero-coupon bonds with maturities 0.5, 1, 1.5, 2, 2.5, 4, 5, 6, 7, 8, 9, 10, 15, 20, 25 and 30 years are used. The one-year bond is essentially a risk-free investment over a one-year investment period. The corresponding zero rates or yields for these bonds are used as the observed input into the Kalman filter. The estimation procedure is carried out on a rolling window basis over the period January 2004 to March 2014, with the window being moved forward by 30 days each time. Six model variants are estimated: completely and essentially affine versions of the one-, two- and three-factor models. Once the parameters of a model have been estimated, the expected returns and covariance matrix of returns for an investment period of one year are calculated and used as inputs for mean-variance optimisation. This procedure ultimately produces 52 estimates of parameters, mean vectors and covariance matrices of yearly zero-coupon bond returns. Table 4.1 provides the mean values of the speed of mean reversion parameter estimates over all the estimation periods. It can be seen that at least one of the factors in each model variant possesses noteworthy mean reversion. Korn and Koziol (2006) highlight that mean reversion plays an important role in the mean-variance framework because it implies that expected returns can be predicted to an extent due to state dependence of the factors.

Tab. 4.1: Mean Speed of Mean Reversion Parameter Estimates

	1-Factor Model	2-Factor Model	3-Factor Model
Completely Affine Specification			
κ_1	0.22266	0.08263	0.47020
κ_2	-	0.43130	0.33242
κ_3	-	-	0.08076
Essentially Affine Specification			
κ_1	0.24305	0.03748	0.00841
κ_2	-	0.40840	0.11906
κ_3	-	-	0.54730

4.2 Mean-Variance Optimisation

The investment set consists of two, three, six or 15 bonds. The two-bond set includes 1- and 15-year bonds, the three-bond set 1-, 5- and 15-year bonds, the six-bond set 1, 2, 5, 10, 15 and 25-year bonds, and the 15-bond set 1-, 1.5-, 2-, 2.5-, 4-, 5-, 6-, 7-, 8-, 9-, 10-, 15-, 20-, 25- and 30-year bonds. Once the expected returns and covariance matrix of returns have been obtained for a given model variant, they are used to determine portfolios which minimise volatility for an expected return of 8%. Only long positions are permitted. In the cases where an expected return of 8% cannot be achieved, the holdings which provide the closest possible expected return to 8% are used.

Tables 4.3 and 4.4 provide summary statistics of the performance of the mean-variance optimised portfolios determined using completely affine and essentially affine model variants, respectively. This refers to the performance realised over the one-year period immediately succeeding the estimation period used to determine the model parameter values. The statistics summarise the performance of 52 strategies over the one-year investment periods for which they were optimised. Note that there is an 11-month overlap in each of the strategies which will affect the quoted statistics. In particular, Tables 4.3 and 4.4 show the mean realised portfolio return, standard deviation of realised portfolio returns, realised Sharpe Ratio, mean realised excess return above the expected return, proportion of instances the realised return exceeds the expected return, and the mean excess return associated with these instances (indicated by *). Expected return refers to the target return of 8%, except in cases where this target could not be obtained. In these cases, the

expected return refers to that given by the relevant optimised portfolio holdings.

Benchmark portfolios for each set of bonds are constructed for comparison purposes in order to ascertain whether the estimation and optimisation procedures are worthwhile. The observed yields to maturity of the relevant bonds on the day immediately preceding the one-year investment period in question are used as proxies for the expected returns over this period. Naive benchmark strategies are determined using MATLAB's solver, in conjunction with a long-only holding constraint, to find portfolio weights which provide an expected return of 8%. This is done to make the benchmark portfolios comparable to the portfolios which are mean-variance optimised given a target expected return of 8%. The average performance of these benchmark portfolios over the 52 investment periods is summarised in Table 4.2.

Tab. 4.2: Performance of benchmark portfolios

	2 Bonds	3 Bonds	6 Bonds	15 Bonds
Mean Return	0.07755	0.07984	0.10781	0.12533
Std Dev of Returns	0.10956	0.09604	0.18094	0.21087
Sharpe Ratio	0.18126	0.23070	0.27700	0.32078

Tab. 4.3: Out-of-sample performance of bond portfolios (completely affine)

2 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.06961	0.07220	0.06290
Std Dev of Realised Returns	0.06328	0.10306	0.08754
Realised Sharpe Ratio	0.18841	0.14084	0.05954
Mean Excess Return	-0.01235	-0.00460	-0.01060
Outperformance Ratio	0.32692	0.46154	0.34615
Mean Excess Return*	0.03512	0.05849	0.04304
3 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.07334	0.07995	0.07607
Std Dev of Realised Returns	0.06622	0.10392	0.09327
Realised Sharpe Ratio	0.23636	0.21423	0.19715
Mean Excess Return	-0.00713	0.00432	0.00306
Outperformance Ratio	0.36538	0.59615	0.51923
Mean Excess Return*	0.04510	0.06006	0.05477
6 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.08225	0.08235	0.10730
Std Dev of Realised Returns	0.05810	0.09503	0.16702
Realised Sharpe Ratio	0.42271	0.25955	0.29703
Mean Excess Return	0.00184	-0.00054	0.03278
Outperformance Ratio	0.40385	0.51923	0.55769
Mean Excess Return*	0.04617	0.05891	0.11454
15 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.06012	0.08302	0.11860
Std Dev of Realised Returns	0.08659	0.09904	0.21686
Realised Sharpe Ratio	0.39484	0.25579	0.28091
Mean Excess Return	0.01209	0.00017	0.04412
Outperformance Ratio	0.46154	0.46154	0.51923
Mean Excess Return*	0.06012	0.06513	0.15003

Tab. 4.4: Out-of-sample performance of bond portfolios (essentially affine)

2 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.07802	0.06905	0.07147
Std Dev of Realised Returns	0.06719	0.10400	0.05694
Realised Sharpe Ratio	0.30266	0.10929	0.24210
Mean Excess Return	-0.00040	-0.00931	-0.00235
Outperformance Ratio	0.34615	0.46154	0.24210
Mean Excess Return*	0.04910	0.05416	0.02621
3 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.08120	0.07805	0.08166
Std Dev of Realised Returns	0.06549	0.11004	0.06077
Realised Sharpe Ratio	0.35901	0.18505	0.39448
Mean Excess Return	0.00325	0.00157	0.00853
Outperformance Ratio	0.44231	0.57692	0.44231
Mean Excess Return*	0.04341	0.06358	0.04697
6 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.08411	0.09850	0.09919
Std Dev of Realised Returns	0.07203	0.10229	0.14801
Realised Sharpe Ratio	0.36684	0.39895	0.28041
Mean Excess Return	0.00685	0.01573	0.02659
Outperformance Ratio	0.50000	0.57692	0.48077
Mean Excess Return*	0.04691	0.07098	0.09816
15 Bonds			
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
Mean Realised Return	0.08661	0.09570	0.12074
Std Dev of Realised Returns	0.07097	0.09123	0.20291
Realised Sharpe Ratio	0.40755	0.41671	0.31076
Mean Excess Return	0.00937	0.01228	0.04835
Outperformance Ratio	0.53846	0.55769	0.48077
Mean Excess Return*	0.04695	0.06317	0.14768

A superior mean-variance optimised portfolio is that which provides a higher realised Sharpe ratio. Outperformance refers to superior performance on a risk-adjusted basis, unless otherwise stated. Using a completely affine specification, from Table 4.3, a one-factor model provides significantly superior mean-variance

portfolios for all sets of bonds, relative to its two- and three-factor counterparts. Using an essentially affine specification, from Table 4.4, a one-factor model provides superior mean-variance portfolios consisting of two bonds on average and a three- factor model provides superior three-bond portfolios. A two-factor model is superior when optimising over a set of six and 15 bonds. Therefore, no relationship between the number of factors used in the model and the number of bonds in the portfolio is discernible, unlike the findings of Korn and Koziol (2006). They observe that favourable risk-return profiles can be obtained if there are at least as many risky bonds in the portfolio as there are risk factors in the model.

A completely affine one-factor model only outperforms all essentially affine variants in the case of six bonds. An essentially affine model variant provides superior strategies on average when optimising over sets of two, three and 15 bonds. Essentially affine models also achieve higher mean realised returns in general. The realised returns of the optimised strategies are mostly close to what is expected during periods of relative economic stability; except when a completely affine three-factor model is used to optimise over sets of two, three and six bonds. In these cases, the completely affine three-factor model displays inadequate and erratic forecasting of expected returns as illustrated in Figure 4.1. Figures 4.1 and 4.2 contain plots showing a comparison of realised returns and expected returns for a portfolio of six bonds over the 52 investment periods for completely affine model variants and essentially affine model variants, respectively. Figures 4.3 and 4.4 provide the same for a portfolio of 15 bonds. The instances where there are significant deviations for all model variants across all portfolios coincide with periods of great economic instability. It can be seen that realised returns significantly exceed expected returns during 2008 for all model variants. This observation can be explained by the 2008 global financial crisis, where government bonds became significantly more attractive to investors and, in particular, those of emerging markets. During 2013, however, significant underperformance of bond strategies relative to expectations is observed. The onset of the emerging markets downslide, largely triggered by uncertainty over US Federal Reserve monetary policy and slowing growth in China, serves as an explanation for this poor performance. The essentially affine two- and three-factor models show significantly superior forecasting ability, in comparison to their completely affine equivalents, during these periods of extreme downside volatility.

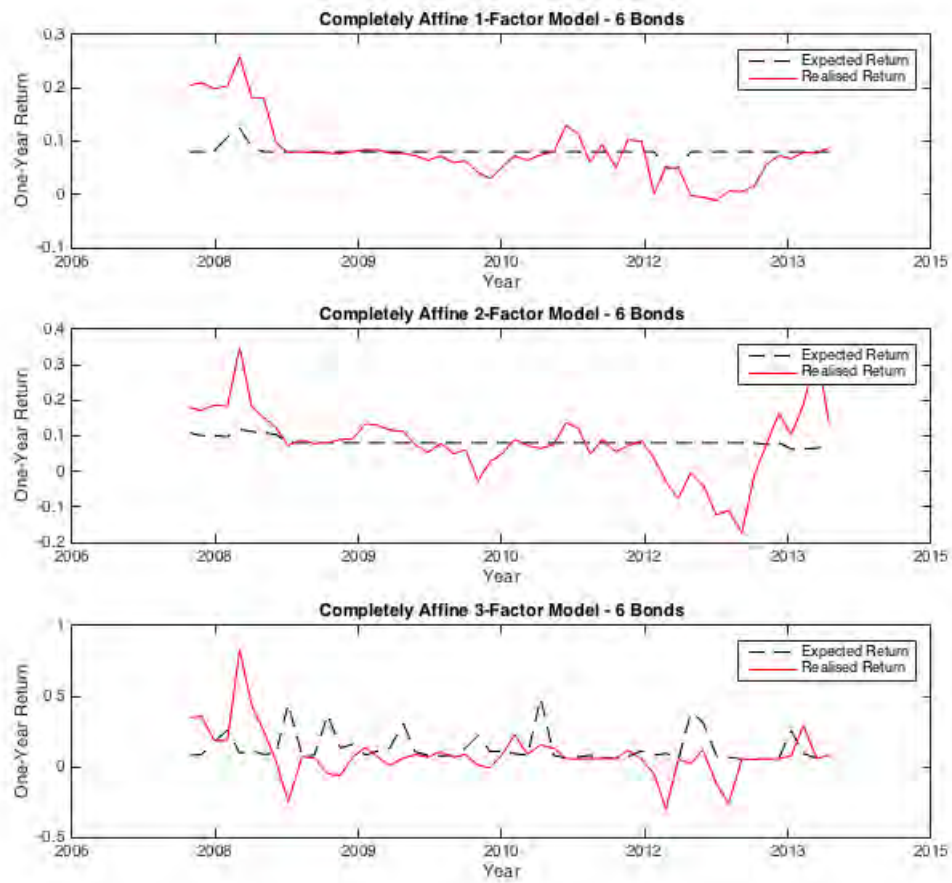**Fig. 4.1:** Expected versus realised returns for a completely affine model applied to a portfolio of six bonds

Fig. 4.2: Expected versus realised returns for an essentially affine model applied to a portfolio of six bonds
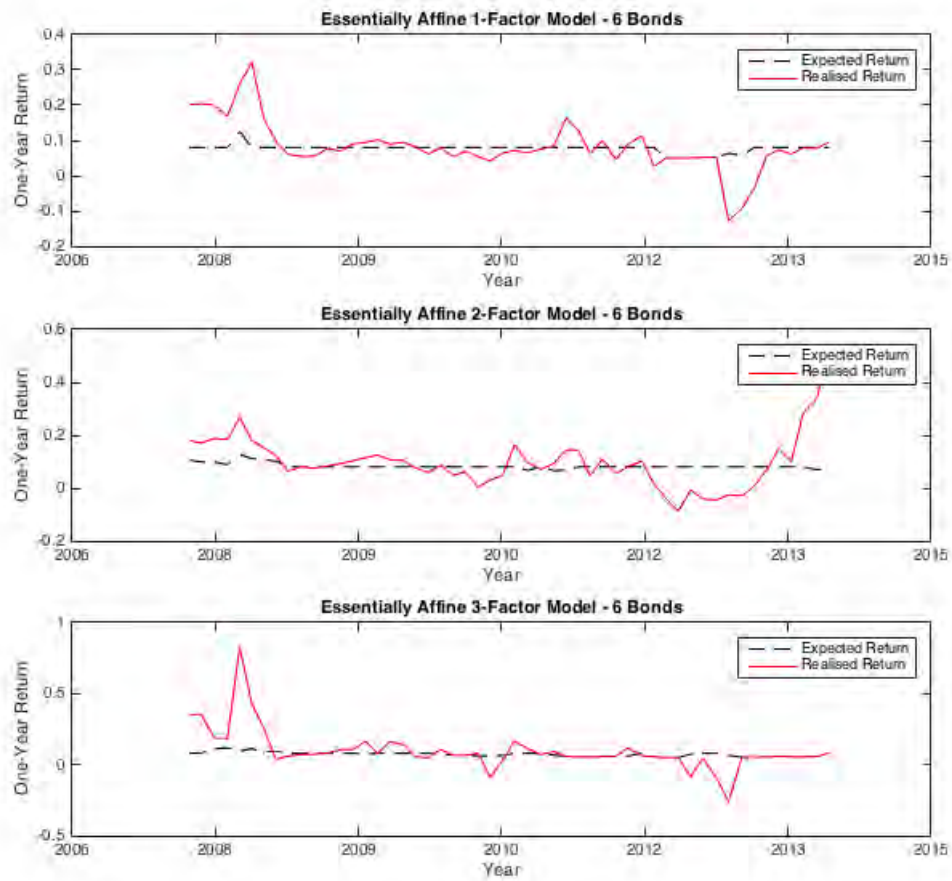
Fig. 4.3: Expected versus realised returns for a completely affine model applied to a portfolio of 15 bonds
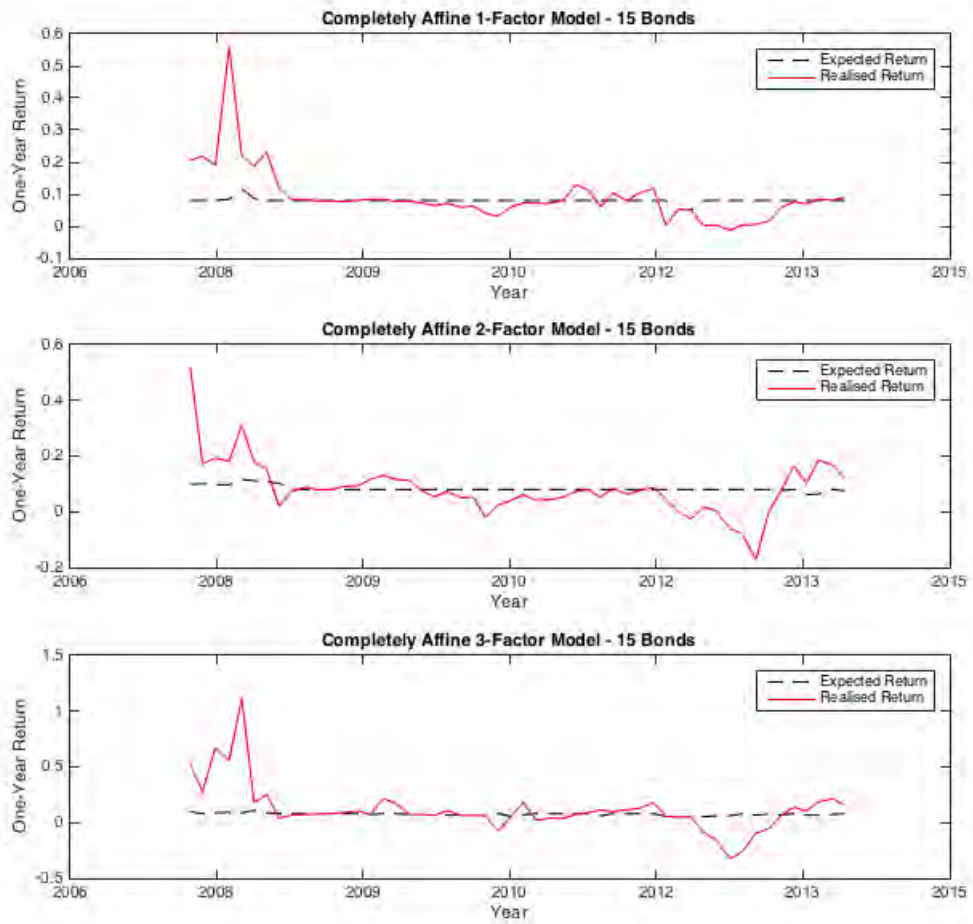
Fig. 4.4: Expected versus realised returns for an essentially affine model applied to a portfolio of 15 bonds
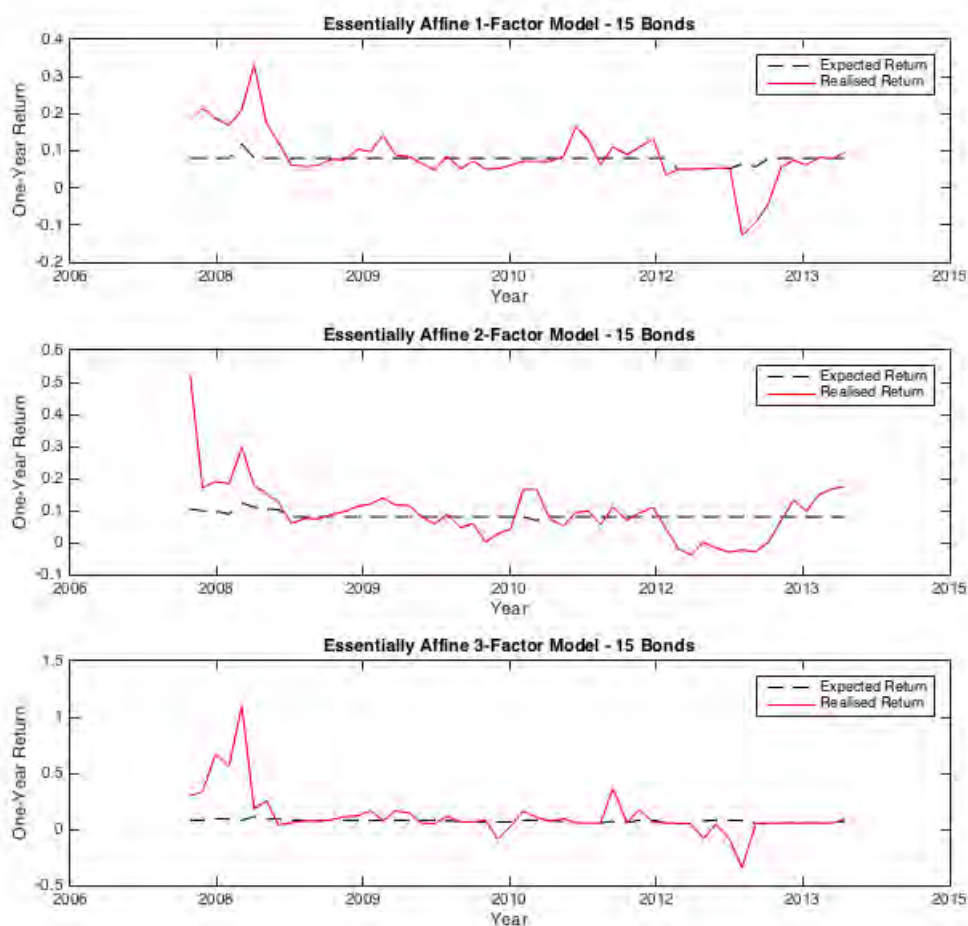

Figure 4.5 shows that the completely affine speed of mean reversion parameter estimates fluctuate sharply over time as opposed to the essentially affine parameters, see Figure 4.6. As aforementioned, a considerable degree of mean reversion makes model estimation worthwhile because this implies some level of predictability of expected returns (Korn and Koziol, 2006). This parameter estimation instability is a possible explanation for the inferior risk-adjusted performance of completely affine models in general —particularly relating to the prediction of expected returns.

Fig. 4.5: Completely affine speed of mean reversion parameter estimates
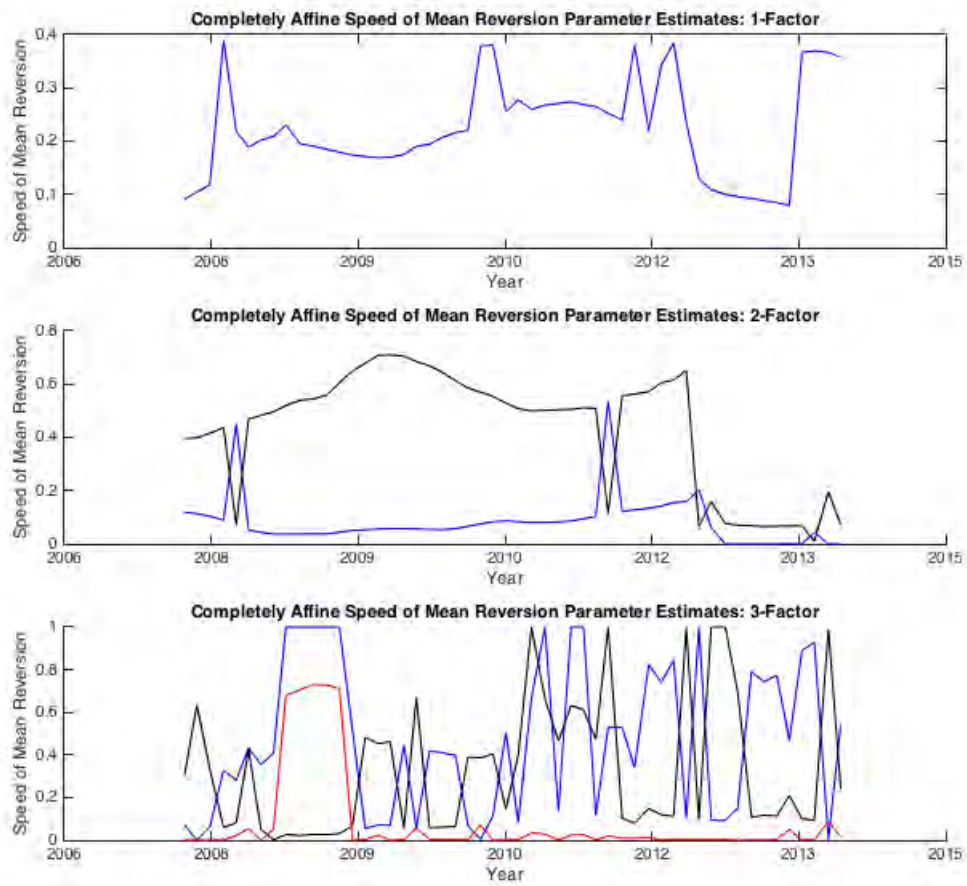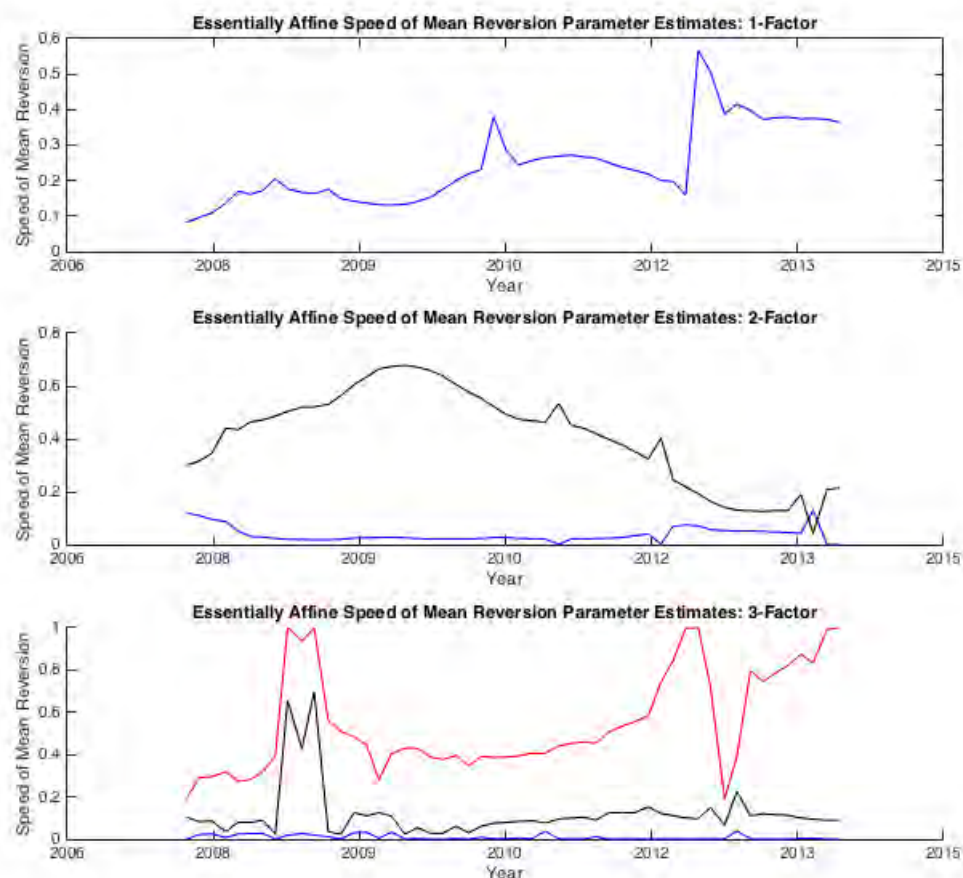
Fig. 4.6: Essentially affine speed of mean reversion parameter estimates

Essentially affine models are not without flaws. They consistently underestimate portfolio volatility relative to the observed standard deviation of realised portfolio returns. Duffee (2002) finds that there is a trade-off between capturing time variation in the conditional variances of yields and time variation in the market price of risk. Therefore, any improvement in forecasting future yields comes at the expense of estimating interest rate volatility. The breaking of the proportionality between the price of risk and interest rate volatility seems to make it more difficult to disentangle the two under the real-world measure using time-series information. Despite this downside of essentially affine models, which is supported in the literature, essentially affine models outperform completely affine models on a risk-adjusted basis. In addition, advanced knowledge of the underestimation of volatility mitigates this downside.

Table 4.2 shows that the mean-variance optimised portfolios provide signifi-

cantly superior risk-adjusted performance than the benchmark portfolios. Therefore, the application of the mean-variance framework in the optimisation of bond portfolios has its merits in that it facilitates the attainment of more attractive risk-return profiles.

Chapter 5

Conclusion

As documented in Duffee and Stanton (2012), the risk premia of the essentially affine model are estimated with strong negative bias and a high variance. This is because these parameters are identified only under the real-world measure. The parameters which are identified under both the risk-neutral and real-world measures —the speed of mean reversion, the long-term mean level and local volatility —are estimated with very little bias. There is difficulty in precisely identifying the level of mean reversion parameters associated with each of the factors in a multifactor setting. In addition, the volatility parameters of an essentially affine multifactor model are estimated with less precision, usually underestimated. This observation is supported in the literature —Duffee (2002) finds that there is a trade-off between forecasting future yields and estimating interest rate volatility.

However, despite these parameter estimation issues, essentially affine models, when used in a mean-variance optimisation setting, still largely outperform completely affine models on a risk-adjusted basis. Essentially affine two- and three-factor models demonstrate superior ability in forecasting expected returns, especially in times of extreme downside volatility. The predictive capability of a model is closely tied to the speed of mean reversion —a high degree of mean reversion implies stronger predictive capability. At least one factor in each of the model variants possesses considerable mean reversion, indicating that the estimation procedure is worthwhile.

However, the estimates of completely affine speed of mean reversion parameters fluctuate sharply over time, which is a cause for concern for reliable prediction of expected returns and a possible explanation for inferior risk-adjusted performance. There is no discernible relationship between the number of bonds included in the portfolio and the number of factors used in the model that enhances portfolio performance. Therefore, the results of this dissertation do not confirm the finding of Korn and Koziol (2006) that favourable risk-return profiles can be obtained if there are at least as many risky bonds in the portfolio as there are risk factors used

in the model.

Finally, the mean-variance optimised portfolios outperform the benchmark portfolios on a risk-adjusted basis, which justifies the estimation and optimisation procedures. The more flexible market price of risk specification of essentially affine models enables superior portfolio performance through improved prediction of expected returns. The trade-off for this increased flexibility, however, is the consistent underestimation of the volatility of returns of the optimised portfolio strategies. That being said, there is potential for mitigation of any detrimental effects of this volatility underestimation given knowledge of it in advance. Therefore, the parameter estimation issues associated with essentially affine models are outweighed by the benefits of superior returns forecasting.

Potential future research avenues include the application of essentially affine models within the duration-constrained mean-variance framework of Caldeira *et al.* (2016), as well as the link between essentially affine mean-variance optimised portfolios and the bond portfolio strategies explored by Puhle (2008).

Bibliography

Babbs, S. H. and Nowman, K. B. (1999). Kalman filtering of generalized vasicek term structure models, *Journal of Financial and Quantitative Analysis* **34**(01): 115–130.

Bolder, D. J. (2001). Affine term-structure models: Theory and implementation, *Available at SSRN 1082826* .

Brennan, M. J. and Schwartz, E. S. (1980). Conditional predictions of bond prices and returns, *The Journal of Finance* **35**(2): 405–417.

Brigo, D. and Mercurio, F. (2007). *Interest rate models-theory and practice: with smile, inflation and credit*, Springer Science & Business Media.

Caldeira, J. F., Moura, G. V. and Santos, A. (2012). Bond portfolio optimization: a dynamic heteroskedastic factor model approach, *Technical report*, Working Paper, Department of Economics, Universidade Federal de Santa Catarina.

Caldeira, J. F., Moura, G. V. and Santos, A. A. (2016). Bond portfolio optimization using dynamic factor models, *Journal of Empirical Finance* **37**: 128–158.

Campbell, J. Y. and Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view, *The Review of Economic Studies* **58**(3): 495–514.

Cox, J. C., Ingersoll Jr, J. E. and Ross, S. A. (1985). A theory of the term structure of interest rates, *Econometrica: Journal of the Econometric Society* pp. 385–407.

Dai, Q. and Singleton, K. (2000). Specification analysis of term structure of interest rates, *Journal of Finance* **55**: 1943–78.

De Jong, F. (1999). Time-series and cross-section information in affine term structure models, *Technical report*, CEPR Discussion Papers.

Duan, J.-C. and Simonato, J.-G. (1999). Estimating and testing exponential-affine term structure models by kalman filter, *Review of Quantitative Finance and Accounting* **13**(2): 111–135.

Duarte, J. (2004). Evaluating an alternative risk preference in affine term structure models, *Review of Financial Studies* **17**(2): 379–404.

Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models, *The Journal of Finance* **57**(1): 405–443.

Duffee, G. R. and Stanton, R. H. (2012). Estimation of dynamic term structure models, *The Quarterly Journal of Finance* **2**(02): 1250008.

Duffie, D. and Kan, R. (1996). A yield-factor model of interest rates, *Mathematical finance* **6**: 379–406.

Fisher, M. and Gilles, C. (1996). Estimating exponential-affine models of the term structure, *Unpublished working paper, Federal Reserve Board, Washington, DC* .

Geyer, A. L. and Pichler, S. (1999). A state-space approach to estimate and test multifactor cox-ingersoll-ross models of the term structure, *Journal of Financial Research* **22**(1): 107–130.

Hull, J. C. and White, A. D. (1994). Numerical procedures for implementing term structure models i: Single-factor models, *The Journal of Derivatives* **2**(1): 7–16.

Korn, O. and Koziol, C. (2006). Bond portfolio optimization: A risk-return approach, *Technical report*, CFR Working Paper.

Lund, J. (1997). Non-linear kalman filtering techniques for term-structure models, *Work* .

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

Puhle, M. (2008). Static bond portfolio optimization, *Bond Portfolio Optimization* pp. 41–84.

Vasicek, O. (1977). An equilibrium characterization of the term structure, *Journal of financial economics* **5**(2): 177–188.

Wilhelm, J. (1992). Fristigkeitsstruktur und zins anderungsrisiko - voru berlegungen zu einer markowitz-theorie des bond-portfolio-management, *Zeitschrift fu r betriebswirtschaftliche Forschung* **44**: 209–246.

Appendix A

A.1 Completely Affine Model Parameter Estimates

The completely affine model parameters determined using maximum likelihood estimation by means of the Kalman filter are displayed in Table A.1.

Tab. A.1: Mean Parameter Estimates for Completely Affine Model Variants

	Completely Affine Specification		
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
θ_1	0.12521	0.08711	0.45465
θ_2	-	0.08737	0.46236
θ_3	-	-	0.45770
κ_1	0.22266	0.08263	0.47020
κ_2	-	0.43130	0.33242
κ_3	-	-	0.08076
σ_1	0.06482	0.03466	0.03117
σ_2	-	0.02361	0.03388
σ_3	-	-	0.01219
λ_1	-0.01544	-0.00761	-0.00247
λ_2	-	0.00501	-0.00575
λ_3	-	-	0.00934
σ_k	0.00432	0.00182	0.00120

A.2 Essentially Affine Model Parameter Estimates

The essentially affine model parameters determined using maximum likelihood estimation by means of the Kalman filter are displayed in Table A.2.

Tab. A.2: Mean Parameter Estimates for Essentially Affine Model Variants

	Essentially Affine Specification		
	1-Factor Model	**2-Factor Model**	**3-Factor Model**
θ_1	0.09696	0.01058	0.44772
θ_2	-	0.00977	0.45571
θ_3	-	-	0.39300
κ_1	0.24305	0.03748	0.00841
κ_2	-	0.40840	0.11906
κ_3	-	-	0.54730
σ_1	0.02892	0.00248	0.00111
σ_2	-	0.00085	0.04689
σ_3	-	-	0.00156
λ_{11}	0.02710	0.05547	-0.13548
λ_{12}	-0.31069	-0.65814	-0.24653
λ_{21}	-	-0.00513	0.22887
λ_{22}	-	0.39951	-0.51190
λ_{31}	-	-	-0.14264
λ_{32}	-	-	0.32336
σ_k	0.00473	0.00201	0.00118