Optimal bond portfolios with fixed time to maturity

Patrik Andersson, Stockholm University*
Andreas N. Lagerås, Stockholm University and AFA Insurance†

Abstract

We study interest rate models where the term structure is given by an affine relation and in particular where the driving stochastic processes are so-called generalised Ornstein-Uhlenbeck processes.

For many institutional investors it is natural to consider investment in bonds where the time to maturity of the bonds in the portfolio is kept fixed over time. We show that the return and variance of such a portfolio of bonds which are continuously rolled over, also called rolling horizon bonds, can be expressed using the cumulant generating functions of the background driving Lévy processes associated with the OU processes. This allows us to calculate the efficient mean-variance portfolio. We exemplify the results by a case study on euro swap rates.

We also show that if the short rate, in a risk-neutral setting, is given by a linear combination of generalised OU processes, the implied term structure can be expressed in terms of the cumulant generating functions. This makes it possible to quite easily see what kind of term structures can be generated with a particular short rate dynamics.

Keywords: Interest rate models, rolling horizon bonds, generalised Ornstein-Uhlenbeck processes, affine term structure, mean variance portfolio

*Department of Mathematics, Stockholm University, SE-106 91 Stockholm, Sweden. *E-mail*: patrika@math.su.se. *To whom correspondence should be addressed.*

†AFA Insurance, SE-106 27 Stockholm, Sweden. *E-mail*: andreas.lageras@afaforsakring.se. The opinions expressed in this paper are not necessarily those of AFA Insurance.

1 Introduction

Bonds are an important part of many investors' portfolios. They are sometimes used as a complement to equities since they in general have lower risk than equities and since they give diversification. Some institutional investors such as insurance companies are more or less mandated to hold bonds in order to match the interest rate sensitivity of their liabilities. A theory of bond portfolios is therefore quite necessary to many practitioners. An essential difference between bonds and equities is that the former have fixed times of maturity, whereas the latter do not.

This has consequences even in a one-period portfolio model. Consider a default free bond without coupon payments, a so-called zero coupon bond. If the investment horizon equals the time of maturity, the bond is risk free, i.e. the return over the period is known with certainty. If the horizon is longer than the time to maturity, we know what payment we will receive from the bond, but what return one may get from reinvesting this payment until the horizon may be uncertain. If the time period is shorter than the time to maturity, the price of the zero coupon bond at the horizon is uncertain, but since we know that the price must equal the face value at the time of maturity, we have a "pull-to-par" effect, which gives the price a drift toward the face value.

Contrast this with the usual random-walk models for stock prices (or log-prices), where the uncertainty grows with the length of the investment period, say with the square root of time if one measures uncertainty with standard deviation of price (or log price).

One way to get around this difference is to consider not an investment in a single bond with a given time to maturity, but rather an investment in a strategy where bonds with different times of maturity are repeatedly sold and bought so that the time to maturity, i.e. the duration, is more or less fixed. [Rutkowski, 1999] introduced these artificial securities as 'rolling horizon bonds" and [Ekeland and Taflin, 2005] introduced these as "roll-overs" in a framework for bond portfolios. One particular way of thinking of a rolling horizon bond with a given duration is that it is like a bank account where one has to give a notice to withdraw the money, which can thereafter only be accessed after a time equalling the duration in length.

The rolling horizon bond, even if an artificial and synthetic security, should not be considered (only) as a mathematical tool to simplify the analysis of bond portfolios, since this is close to how institutional investors can think about their bond holdings. If they hold bonds with respect to certain liabilities, it is reasonable for them to choose a target duration of the assets which is close to that of the liabilities in order to minimise the effect of interest rate changes to the net funding ratio. To be specific, consider a pension fund that has guaranteed a certain future payment. The only risk free way to hedge this liability is to buy a default free bond with the time of maturity equal to the time of payment. If the pension fund is in a steady state, adding new liabilities with the same time to payment, the duration of the liabilities will be stable in time as well, and thus also the target duration of the hedging assets. Deviations from this duration might be appropriate if one expects the return on bonds to be better for another duration, but this decision must be made with respect to the duration of the liabilities, since this adds additional risk. If, on the other hand, the pension fund is in run-off, the duration of the hedging assets does not need to be continuously adjusted, and one may simply wait and let the maturing bonds match the liabilities.

In this paper we restrict the study of portfolios of rolling horizon bonds to models of the yield curve where it can be represented as linear combination of some stochastic processes. One can compare these with factor models for equity returns such as the CAPM and the Fama-French three factor model. Such factor models are useful since they show where the main variation in returns comes from. For yield curves, one has observed that most of their changes in shapes can be represented with parallel shifts, and to a lesser extent with changes in slope and curvature. It is therefore natural to base a yield curve model on these factors.

Since we are interested in actual future return we are mostly concerned with what is sometimes called the objective probability measure. Some synonyms are the data generating measure and the real world measure. This is in contrast with the risk neutral measure which is implied by the prices in an arbitrage free market. We reason about the so called objective measure in order to make investment decisions, and this measure can hardly ever be fully known, so the decision is more or less dependent on subjective opinions. Thus it might just as well be called the subjective measure. We will however also touch upon risk neutral modelling. This comes naturally as we will use generalised Ornstein-Uhlenbeck processes as factors in our model — these are general in the sense that they may be driven by any Lévy process rather than just a Brownian motion. It is well-known that the yield curve can only have such an affine representation and still belong to an arbitrage free model if the factors have a certain dynamic which include the generalised Ornstein-Uhlenbeck processes, see [Duffie et al., 2003]. The lack of arbitrage also puts restraints on how the loadings on the factors depend on the time to maturity of the interest rates. We will show how these loadings look in a risk neutral affine model driven by a generalised Ornstein-Uhlenbeck process, which might be of independent interest.

However, we will not in general restrict ourselves to models that are arbitrage free. The reason is that there are models which are known to represent both the shape and dynamics of the yield curve without being arbitrage free. It is in general also hard to estimate the market price of risk related to the different factors, and this is needed in order to transition between the risk neutral and the objective probability measure. We exemplify some results with the dynamic Nelson-Siegel model, which has been used, e.g., by several central banks. It is not arbitrage free in the mathematical sense but may be practically indistinguishable from arbitrage free, [Coroneo et al., 2008].

The structure of the paper is as follows: In Section 2 we state some basic results regarding interest rates and rolling horizon bonds and show how to express the return of these bonds using forward rates. In Section 3 we recall the definition of an affine term structure and show that the return of a portfolio of rolling horizon bonds in this case has a simple structure. The generalized Ornstein-Uhlenbeck processes are defined in Section 4, together with some results regarding risk neutral modelling. In Section 5 we consider portfolios of rolling horizon bonds and perform a case study using euro swap rates. Section 6 concludes the paper.

The novelties in this paper are the following: In Section 4.2 we show that if the short rate is described by a certain Ornstein-Uhlenbeck process, one can easily judge if this process corresponds to a reasonable yield curve. This is possible since the yield curve can be written in terms of the generating functions. In Section 5.1 we combine the results in [Rutkowski, 1999] and [Barndorff-Nielsen and Shephard, 2003] to model the return of rolling horizon bonds. We derive explicit expressions for the mean and covariance, thus making a mean-variance optimization possible. In Section 5.2 we suggest a non-parametric method of estimating the unknown parameters in the model. In Section 5.3 we do a case study. The results of the case study are in some sense negative, since the constructed portfolios do not perform very well out of sample. However, it is common that mean-variance portfolios perform poorly out of sample. Or paper thus shows that more elaborate methods needs to be used to achieve betters performance.

2 Interest rates and bonds

Let $\mathcal{Z}_t(\tau)$ be the value at time t of a zero coupon bond with maturity $t+\tau$ and nominal value 1 EUR. Let $f_t(\tau', \tau'')$ be the simple forward rate contracted at time t for the period $t+\tau'$ to $t+\tau''$, and let $f_t(\tau)$ be the instantaneous forward rate at time t for maturity $t+\tau$. For any mesh $0=\tau_0<\cdots<\tau_n=\tau$,

$$\mathcal{Z}_t(\tau)=\prod_{i=1}^{n}(1+(\tau_i-\tau_{i-1})f_t(\tau_{i-1},\tau_i))^{-1}=\exp\left\{-\int_0^{\tau} f_t(u)du\right\}. \qquad (1)$$

Expressed with the continuously compounded spot rate $y_t(\tau)$,

$$\mathcal{Z}_t(\tau)=\exp\{-\tau y_t(\tau)\},$$

so that

$$y_t(\tau)=\frac{1}{\tau}\int_0^{\tau} f_t(u)du. \qquad (2)$$

For future reference, note that equation (1) implies that

$$\frac{\mathcal{Z}_t(\tau')}{\mathcal{Z}_t(\tau'')}=1+(\tau''-\tau')f_t(\tau',\tau'')=\exp\left\{\int_{\tau'}^{\tau''} f_t(u)du\right\}, \qquad (3)$$

for $\tau'<\tau''$.

We assume that the forward curve $f_t(\tau)$ is continuous in τ for all times t, that $f_t(\tau)$ as a function of t is Riemann integrable for all τ and that

$$\lim_{\tau^*\to\tau}\sup_{t>0}|f_t(\tau)-f_t(\tau^*)|=0.$$

Let $R_t^d(\tau)$, d for *discrete*, be the value at time t of the following strategy: Start at time 0 with 1 EUR. At all times $0=t_0<t_1<t_2<\cdots$ rebalance to hold only zero coupon bonds with time to maturity τ. Let $\Delta_i=t_i-t_{i-1}$.

Proposition 1 ([Rutkowski, 1999])**.**

1. $$R_{t_i}^d(\tau)=\frac{\mathcal{Z}_{t_i}(\tau)}{\mathcal{Z}_0(\tau)}\prod_{k=1}^{i}\left(1+\Delta_k f_{t_k}(\tau-\Delta_k,\tau)\right),\quad i=0,1,\ldots$$

2. *As the trading frequency tends to infinity, R_t^d converges to*

$$R_t(\tau)=\frac{\mathcal{Z}_t(\tau)}{\mathcal{Z}_0(\tau)}\exp\left\{\int_0^t f_s(\tau)ds\right\},\quad t\in\mathbb{R}_+.$$

The proof of the first part follows by induction using (3) and the second part follows from the first by straightforward Riemann integration.

The strategy of continuously rebalancing a portfolio to hold bonds with a single time to maturity is said to produce a "rolling horizon bond" by [Rutkowski, 1999] and a "roll-over" by [Ekeland and Taflin, 2005]. If one at any point in time were to stop rolling over and rather let the bond mature it would take τ time units until maturity when the principal of the bond would be paid. In effect one can say that $R_t(\tau)$ describes the value of a bank account where notice of withdrawal must be made τ time units in advance [Ekeland and Taflin, 2005].

The expression for $R_t(\tau)$ can be stated entirely with forward rates using (1):

$$R_t(\tau) = \exp\left\{ -\int_0^\tau (f_t(u) - f_0(u))du + \int_0^t f_s(\tau)ds \right\}. \tag{4}$$

Let $X_t(\tau) = \log R_t(\tau)$, so that

$$X_t(\tau) = -\int_0^\tau (f_t(u) - f_0(u))du + \int_0^t f_s(\tau)ds. \tag{5}$$

Note that $R_t(0)$ is the value of a money market account which accumulates the short rate and is available for immediate withdrawal. Thus, the excess return of rolling bonds with time to maturity τ compared with holding this money market account, has the following, quite symmetric, expression:

$$X_t(\tau) - X_t(0) = -\int_0^\tau (f_t(u) - f_0(u))du + \int_0^t (f_s(\tau) - f_s(0))ds.$$

The average logarithmic return of the rebalancing strategy is

$$\frac{X_t(\tau)}{t} = -\frac{1}{t}\int_0^\tau (f_t(u) - f_0(u))du + \frac{1}{t}\int_0^t f_s(\tau)ds = \frac{1}{t}\int_0^t f_s(\tau)ds + O\left(t^{-1}\right).$$

Thus if the forward rate process $f_t(\tau)$ is ergodic with stationary distribution F, the average logarithmic return from a rolling horizon bond will converge almost surely to $\mathbb{E}[f_\infty(\tau)]$, and the average geometric return will converge almost surely to $\exp\{\mathbb{E}[f_\infty(\tau)]\} - 1$, where $f_\infty(\tau) \sim F$.

3 Affine term structure and portfolios

In the rest of this paper we will consider the case where the forward rate is given by an affine process, i.e. the dynamics of the forward curve can be described by an affine relationship

$$f_t(\tau) = \boldsymbol{\kappa}(\tau)'\mathbf{F}_t = (\kappa_1(\tau), \ldots, \kappa_n(\tau))(F_{1t}, \ldots, F_{nt})'. \tag{6}$$

Note that one of the factors F_{kt} may be constant, equal to 1 say, in t.

By (2), this is equivalent to assuming an affine relationship for the spot rates:

$$y_t(\tau) = \frac{1}{\tau}\bar{\boldsymbol{\kappa}}(\tau)'\mathbf{F}_t,$$

where $\bar{\boldsymbol{\kappa}}(\tau) = \int_0^\tau \boldsymbol{\kappa}(u)du$.

One example of an affine model is the Vasiček model, where

$$f_t(\tau) = \mu\frac{1 - e^{-\lambda\tau}}{\lambda} - \frac{\sigma^2}{2}\left(\frac{1 - e^{-\lambda\tau}}{\lambda}\right)^2 + e^{-\lambda\tau}r_t.$$

Here the short rate r_t is the only non-constant factor.

Another example is the Nelson-Siegel model, where

$$f_t(\tau) = \beta_{0t} + e^{-\gamma\tau}\beta_{1t} + \gamma\tau e^{-\gamma\tau}\beta_{2t}.$$

If one wants an affine model to be arbitrage free, the κ-functions cannot be set willy-nilly, since they are determined by the dynamics of $\mathbf{F}_t$ under the pricing, or so-called risk neutral, probability measure. See [Duffie et al., 2003] for an exhaustive treatise on this matter. Of the examples above, the Vasiček model is arbitrage free whereas the Nelson-Siegel model is not, if $\boldsymbol{\beta}_t$ is given by an Itô process, see [Filipović, 1999].

We will in Section 4.2 give some consideration to models where the dynamics of the factors are specified under the pricing measure. We are, however, mostly interested in the returns of bond investments under the real world probability measure, which is also called the objective or data generating measure. It is sometimes possible to specify a market price of risk process such that the forward rate is an affine process under both the pricing and data generating measure. See, e.g., [Piazzesi, 2009] for a nice overview of some possible relations between the two probability measures in order to keep the affine structure under both of them. Nevertheless, we choose to either give examples of processes that are affine under the risk neutral measure or the real world one. One reason to do this is to be able to estimate the functions κ with as few assumptions as possible, e.g. by using principal components analysis on panel data of interest rates with several maturities.

By combining our assumption (6) with (5),

$$X_t(\tau) = -\bar{\boldsymbol{\kappa}}(\tau)'(\mathbf{F}_t - \mathbf{F}_0) + \boldsymbol{\kappa}(\tau)'\bar{\mathbf{F}}_t,$$

where $\bar{\mathbf{F}}_t = \int_0^t \mathbf{F}_s ds$. One can now consider portfolios constructed by rolling horizon bonds at different maturities on the yield curve. A portfolio that at time zero has a value ν_i invested in a rolling-horizon bond with time to maturity τ_i, and which is never rebalanced between these maturities, will at time t have the value

$$\Pi_t^0 = \sum_i \nu_i R_t(\tau_i) = \sum_i \nu_i \exp\{-\bar{\boldsymbol{\kappa}}(\tau_i)'(\mathbf{F}_t - \mathbf{F}_0) + \boldsymbol{\kappa}(\tau_i)'\bar{\mathbf{F}}_t\}.$$

Note that all ν_i need not be positive if there is a possibility to short the bonds.

4 Ornstein-Uhlenbeck-type models

In this section we recall some well-known facts about Ornstein-Uhlenbeck processes. We use Ornstein-Uhlenbeck processes for modelling the factors in the affine expression of interest rates since they are relatively simple processes that are mean reverting. This is a desirable property since it is unreasonable to assume that interest rates could diverge toward infinity, at least in a non-hyperinflationary economy. We stress again that there are differences between modelling the dynamics of the factors under the risk neutral measure and the objective measure. The risk neutral dynamics should be such that the observed yield curve at any point in time is consistent with the risk neutral pricing. The objective dynamics, on the other hand, describes how the yield curve actually changes in time.

4.1 Generalized Ornstein-Uhlenbeck processes

This section is heavily based on [Barndorff-Nielsen and Shephard, 2003]. We recall the material here for completeness.

Let Z_t be a Lévy process and

$$Y_t \equiv e^{-\lambda t} Y_0 + \int_0^t e^{-\lambda(t-s)} dZ_s. \tag{7}$$

We call Y_t a generalised Ornstein-Uhlenbeck (OU) process with background driving Lévy process (BDLP) Z_t. This process specialises to the ordinary Ornstein-Uhlenbeck process when Z_t is a Brownian motion.

If $\mathbb{E}[\log(1+|Z_1|)] < \infty$, then $Y_t \xrightarrow{d} Y_\infty$ as $t \to \infty$, where $Y_\infty \equiv \int_0^\infty e^{-\lambda t} dZ_t$. If one lets $Y_0 \stackrel{d}{=} Y_\infty$, the process becomes stationary. However, we will mostly assume that Y_0 is constant.

We define the integrated process

$$\bar{Y}_t \equiv \int_0^t Y_s ds = \frac{1-e^{-\lambda t}}{\lambda} Y_0 + \frac{1}{\lambda} \int_0^t (1 - e^{-\lambda(t-u)}) dZ_u.$$

With $\varepsilon(t;\lambda) \equiv (1-e^{-\lambda t})/\lambda$, for $\lambda > 0$, and $\varepsilon(t;0) \equiv t$

$$\bar{Y}_t = \varepsilon(t;\lambda) Y_0 + \int_0^t \varepsilon(t-u;\lambda) dZ_u.$$

We will use the cumulant generating function $l(\theta \ddagger X) \equiv \log \mathbb{E}[e^{\theta X}]$ as a tool for deriving some useful relations between the distributions of Z, Y and $\bar{Y}$. We will mostly consider positive random variables, for which this cumulant function is well defined when $\theta \le 0$, and Gaussian random variables. In the general case one can use a cumulant function based on the characteristic function rather than our preferred one based on the moment generating function transform.

Let

$$\begin{aligned} \zeta(\theta) &\equiv l(\theta \ddagger Z_1), \\ \upsilon_t(\theta) &\equiv l(\theta \ddagger Y_t), \\ \upsilon(\theta) &\equiv l(\theta \ddagger Y_\infty), \\ \upsilon_t^*(\theta) &\equiv l(\theta \ddagger \bar{Y}_t), \end{aligned}$$

where we assume that Y_0 is constant. We also define the joint cumulant generating function $\upsilon_t^\circ(\theta_1,\theta_2) \equiv \log \mathbb{E}[e^{\theta_1 Y_t + \theta_2 \bar{Y}_t}]$. These generating functions are related as follows.

$$\begin{aligned} \zeta(\theta) &= \lambda \theta \upsilon'(\theta), \\ \upsilon_t(\theta) &= \theta e^{-\lambda t} Y_0 + \int_0^t \zeta(\theta e^{-\lambda s}) ds, \\ \upsilon(\theta) &= l\left(\theta \ddagger \int_0^\infty e^{-\lambda t} dZ_t\right) = \int_0^\infty \zeta(\theta e^{-\lambda t}) dt, \\ \upsilon_t^*(\theta) &= \theta \varepsilon(t;\lambda) Y_0 + \int_0^t \zeta(\theta \varepsilon(s;\lambda)) ds, \\ \upsilon_t^\circ(\theta_1,\theta_2) &= \theta_1 e^{-\lambda t} Y_0 + \theta_2 \varepsilon(t;\lambda) Y_0 + \int_0^t \zeta\big(\theta_1 e^{-\lambda s} + \theta_2 \varepsilon(s;\lambda)\big) ds. \end{aligned}$$

4.2 Risk neutral models

When the short rate is a linear combination of independent OU processes under the risk neutral measure, the yield curve is easily expressed with these cumulant generating functions. In order to make the presentation clearer, we start with a one-factor model, i.e.

$$r_t = y_t(0) = f_t(0) = e^{-\lambda t} r_0 + \int_0^t e^{-\lambda(t-s)} dZ_s$$

with Z_t a Lévy process as before. Note that this is the Vasiček model if $Z_t = \mu t + \sigma B_t$, with B_t a standard Brownian motion. When Z_t is a subordinator, i.e. a positive Lévy process, the short rate r_t is positive as well.

Letting $\mathbb{E}[\cdot]$ denote the expectation under the risk neutral measure, we have

$$Z_0(\tau) = e^{-\tau y_0(\tau)} = e^{-\int_0^\tau f_0(u)du} = \mathbb{E}[e^{-\int_0^\tau r_s ds}|r_0] = \mathbb{E}[e^{-\bar{r}_\tau}|r_0] = e^{\upsilon_\tau^*(-1)}.$$

The yield curve is therefore given by

$$y_0(\tau) = -\frac{\upsilon_\tau^*(-1)}{\tau} = \frac{\varepsilon(\tau;\lambda)}{\tau} r_0 - \frac{1}{\tau}\int_0^\tau \zeta\left(-\varepsilon(u;\lambda)\right) du,$$

and the forward curve is

$$f_0(\tau) = -\frac{d}{d\tau}\upsilon_\tau^*(-1) = e^{-\lambda\tau} r_0 - \zeta\left(-\varepsilon(\tau;\lambda)\right).$$

Note that the effect of the current short rate on the yield (forward) curve, given by the coefficient $e^{-\lambda\tau}$, is the same, regardless of what BDLP Z is used in the model. Since both functions $e^{-\lambda\tau}$ and $\varepsilon(\tau;\lambda)/\tau$ are monotone in τ, these models cannot reproduce the empirically observed shocks to the yield curve's curvature, i.e. the 'twist' or 'butterfly' factor.

In the Vasiček model $\zeta(\theta) = \mu\theta + \sigma^2\theta^2/2$, which gives the well known expression

$$f_0(\tau) = e^{-\lambda\tau} r_0 + \mu\frac{1-e^{-\lambda\tau}}{\lambda} - \frac{\sigma^2}{2}\left(\frac{1-e^{-\lambda\tau}}{\lambda}\right)^2.$$

If $r_t = \sum_k X_{k,t}$ where $X_{1,t}, X_{2,t}, \ldots$ are independent but not necessarily identically distributed Ornstein-Uhlenbeck processes, the above calculations are straightforward to generalise, producing

$$f_0(\tau) = \sum_k e^{-\lambda_k\tau} x_{k,0} - \sum_k \zeta_k\left(-\varepsilon(\tau;\lambda_k)\right).$$

In [Barndorff-Nielsen and Shephard, 2003], a generalised OU process with $Y_\infty \sim D$ is called D-OU, and on the other hand, if $Z_1 \sim D$ the process is called OU-D. We reproduce their table of generating functions for some different OU processes, with some change in notation. This can be used to ascertain whether a given OU process has a ζ which produces a reasonable yield curve.

Model	$\zeta(\theta)$
OU-Normal(μ,σ^2)	$\mu\theta + \sigma^2\theta^2/2$
OU-Poisson(μ)	$\mu(e^\theta - 1)$
OU-Gamma(ν,α)	$-\nu \log(1 - \theta/\alpha)$
OU-Inverse Gaussian(δ,γ)	$\delta\big(\gamma - \sqrt{\gamma^2 - 2\theta}\big)$
Normal(μ,σ^2)-OU	$\lambda\mu\theta + \lambda\sigma^2\theta^2$
Gamma(ν,α)-OU	$\lambda\nu\theta/(\alpha - \theta)$
Inverse Gaussian(δ,γ)-OU	$\lambda\delta\theta/\sqrt{\gamma^2 - 2\theta}$

If $\lambda = 0$, then $Y_t \equiv Z_t$, i.e. a pure random walk without any mean reversion. This is not an uninteresting model since the mean reversion might very well be quite weak over short time horizons. In this case, Y_∞ does not exist unless $Z_t \equiv 0$. Here $f_0(\tau) = r_0 - \zeta(-\tau)$, which may make it clearer to understand the effects of the parameters of Z on the shape of the yield curve. If the short rate is a Brownian motion: $r_t = r_0 + \mu t + \sigma B_t$, the forward rate curve is $f_0(\tau) = r_0 + \mu\tau - \sigma^2\tau^2/2$. We immediately see that the drift μ gives the slope, or first derivative, of the forward curve, and the instantaneous variance σ^2 gives the negative of the curvature, or second derivative of the forward curve.

5 Bond portfolios

In this section we consider the problem of forming portfolios of bonds, in particular rolling horizon bonds in a mean-variance framework. We show that under the assumption of an affine term structure driven by generalized Ornstein-Uhlenbeck processes the expectation and covariance of the return of rolling horizon bonds can be given quite explicit expression, thus making a mean-variance optimization tractable. We also provide suggestions for the estimation of the unknown parameters and give an example of an implementation in a case study on euro swap rates.

5.1 Mean-variance portfolio

We would like to find the mean-variance optimal portfolios consisting entirely of rolling-horizon bonds. We will assume that the forward rate for each tradable maturity τ is

$$f_t(\tau) = \boldsymbol{\kappa}(\tau)'\mathbf{F}_t + e_t(\tau),$$

where $\mathbf{F}_t$ is a d-dimensional generalised OU process, under the objective measure, with independent components and where $e_t(\tau)$ is a residual independent of $\mathbf{F}_t$ with stationary distribution. We set the current time to $t = 0$ and will consider a one-period investment with horizon Δt. Using Eq. (5) we get the log-return at the investment horizon of a rolling horizon bond based on zero-coupon bonds of maturity τ

$$X_{\Delta t}(\tau) = -\bar{\boldsymbol{\kappa}}(\tau)'(\mathbf{F}_{\Delta t} - \mathbf{F}_0) + \boldsymbol{\kappa}(\tau)'\bar{\mathbf{F}}_{\Delta t} + \bar{e}_{\Delta t}(\tau),$$

where we set the integrated residual

$$\bar{e}_{\Delta t}(\tau) = -\int_0^\tau (e_{\Delta t}(u) - e_0(u))du + \int_0^{\Delta t} e_s(\tau)ds.$$

The expected return is then

$$\mathbb{E}\left[R_{\Delta t}(\tau)\right] = \mathbb{E}\Big[\exp\big\{-\bar{\boldsymbol{\kappa}}(\tau)'(\mathbf{F}_{\Delta t} - \mathbf{F}_0) + \boldsymbol{\kappa}(\tau)'\bar{\mathbf{F}}_{\Delta t}\big\}\Big]\mathbb{E}\Big[\exp\left\{\bar{e}_{\Delta t}(\tau)\right\}\Big].$$

We see that the expected return is effectively the product of two moment generating functions. Thus when a factor F is given by an OU process Y, the cumulant generating function of quantities such as $-\bar{\eta}(Y_t - Y_0) + \eta\bar{Y}_t$ is of interest. We have that

$$\begin{aligned} l(\theta \ddagger -\bar{\eta}(Y_t - Y_0) + \eta\bar{Y}_t) &= \theta\bar{\eta}Y_0 + l(\theta \ddagger -\bar{\eta}Y_t + \eta\bar{Y}_t) \\ &= \theta\bar{\eta}Y_0 + \upsilon_t^\circ(-\bar{\eta}\theta, \eta\theta) \\ &= \theta(\lambda\bar{\eta} + \eta)\varepsilon(t;\lambda)Y_0 \\ &\quad + \int_0^t \zeta\Big(\theta\big(-\bar{\eta}e^{-\lambda s} + \eta\varepsilon(s;\lambda)\big)\Big)ds. \end{aligned}$$

From the above result we therefore get that

$$\mathbb{E}\left[R_{\Delta t}(\tau)\right] = \exp\left\{\sum_{i=1}^d (\lambda_i\bar{\kappa}_i(\tau) + \kappa_i(\tau))\,\varepsilon(\Delta t;\lambda)F_{i,0} + S_\tau\right\}\mathbb{E}\Big[\exp\left\{\bar{e}_{\Delta t}(\tau)\right\}\Big],$$

where

$$S_\tau = \sum_{i=1}^{d} \int_0^{\Delta t} \zeta_i\big(-\bar{\kappa}_i(\tau)e^{-\lambda_i s} + \kappa_i(\tau)\varepsilon(s;\lambda_i)\big)ds,$$

and similarly

$$\mathrm{Cov}\left(R_{\Delta t}(\tau_1), R_{\Delta t}(\tau_2)\right) = \mathbb{E}\left[R_{\Delta t}(\tau_1)\right]\mathbb{E}\left[R_{\Delta t}(\tau_2)\right]\left[\frac{\mathbb{E}\left[R_{\Delta t}(\tau_1)R_{\Delta t}(\tau_2)\right]}{\mathbb{E}\left[R_{\Delta t}(\tau_1)\right]\mathbb{E}\left[R_{\Delta t}(\tau_2)\right]} - 1\right]$$

$$= \mathbb{E}\left[R_{\Delta t}(\tau_1)\right]\mathbb{E}\left[R_{\Delta t}(\tau_2)\right]\left[\frac{e^{S_{\tau_1,\tau_2}}\mathbb{E}\big[\exp\left\{\bar{e}_{\Delta t}(\tau_1) + \bar{e}_{\Delta t}(\tau_2)\right\}\big]}{e^{S_{\tau_1}+S_{\tau_2}}\mathbb{E}\big[\exp\left\{\bar{e}_{\Delta t}(\tau_1)\right\}\big]\mathbb{E}\big[\exp\{\bar{e}_{\Delta t}(\tau_2)\}\big]} - 1\right],$$

where

$$S_{\tau_1,\tau_2} = \sum_{i=1}^{d} \int_0^{\Delta t} \zeta_i\big(-\left(\bar{\kappa}_i(\tau_1) + \bar{\kappa}_i(\tau_2)\right)e^{-\lambda_i s} + \left(\kappa_i(\tau_1) + \kappa_i(\tau_2)\right)\varepsilon(s;\lambda_i)\big)ds.$$

We note that the (co)variance in the return is a consequence of the non-linearity of the cumulant generating function of the factor processes together with the (possible) dependence between the residuals of different maturities. This is natural since a linear cumulant generating function would imply zero variance.

5.2 Estimation of the mean-variance portfolio

What remains from the previous section is the estimation of the factor processes, i.e. of λ and ζ, together with

$$\mathbb{E}\big[\exp\left\{\bar{e}_{\Delta t}(\tau)\right\}\big] \text{ and } \mathbb{E}\big[\exp\left\{\bar{e}_{\Delta t}(\tau_1) + \bar{e}_{\Delta t}(\tau_2)\right\}\big].$$

The first problem can be handled by choosing an appropriate family of distributions and using a likelihood or least-squares based estimation. We however also propose a non-parametric method based on noting that the cumulant generating function can be written as

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{c_n}{n!}s^n,$$

where c_n is the nth cumulant, provided all cumulants exist. Some calculations then show that

$$
\begin{aligned}
&\int_0^t \zeta\big(-\bar{\kappa}e^{-\lambda s} + \kappa\varepsilon(s;\lambda)\big)ds \\
&= \sum_{n=1}^{\infty} \frac{c_n}{n!}\frac{1}{\lambda^n}\left[\sum_{k=0}^{n}\binom{n}{k}\left(-\lambda\bar{\kappa}-\kappa\right)^{n-k}\kappa^k\varepsilon\left(t,\lambda(n-k)\right)\right] \\
&\equiv \sum_{n=1}^{\infty}\frac{c_n}{n!}a_n(t;\kappa,\bar{\kappa}).
\end{aligned}
$$

The above sum seems to converge quite fast so that by estimating the first few cumulants and truncating the sum we should have a reasonable estimate of the integral.

Toward this, let us assume that we have observed a process Y_t of the type of Eq. (7) at intervals Δ and we would like to find estimates of the cumulants of Z_1. We get that

$$
Y_{t+\Delta} = e^{-\lambda\Delta}Y_t + \int_t^{t+\Delta} e^{-\lambda(t+\Delta-s)}dZ_s \equiv aY_t + \epsilon_\Delta = \mathbb{E}\left[\epsilon_\Delta\right] + aY_t + (\epsilon_\Delta - \mathbb{E}\left[\epsilon_\Delta\right]).
$$

Thus, regressing $Y_{t+\Delta}$ on Y_t we get estimates $\widehat{\mathbb{E}\left[\epsilon_\Delta\right]}$, $\hat{a}$ and $\hat{\epsilon}_\Delta$, and we can set $\hat{\lambda} = -\frac{\log\hat{a}}{\Delta}$. Now,

$$
\begin{aligned}
l(\theta \ddagger \epsilon_\Delta) &= l\left(\theta \ddagger \int_0^\Delta e^{-\lambda s}dZ_s\right) = \int_0^\Delta \zeta(\theta e^{-\lambda s})ds \\
&= \int_0^\Delta \sum_{n=1}^{\infty}\frac{c_n}{n!}\theta^n e^{-\lambda n s}ds = \sum_{n=1}^{\infty}\frac{c_n}{n!}\theta^n\varepsilon(\Delta,\lambda n).
\end{aligned}
$$

We then have as a natural estimator

$$
\hat{c}_n(Z_1) = \frac{\hat{c}_n(\epsilon_\Delta)}{\varepsilon(\Delta,\hat{\lambda}n)},
$$

where we can estimate the cumulants of ϵ_Δ by noting that the nth cumulant can be expressed as a polynomial in the n first moments, e.g.

$$
\begin{aligned}
c_1 =&\mathbb{E}\left[\epsilon_\Delta\right], \\
c_2 =&\mathbb{E}\left[(\epsilon_\Delta - \mathbb{E}\left[\epsilon_\Delta\right])^2\right], \\
c_3 =&\mathbb{E}\left[(\epsilon_\Delta - \mathbb{E}\left[\epsilon_\Delta\right])^3\right], \\
c_4 =&\mathbb{E}\left[(\epsilon_\Delta - \mathbb{E}\left[\epsilon_\Delta\right])^4\right] - 3\left(\mathbb{E}\left[(\epsilon_\Delta - \mathbb{E}\left[\epsilon_\Delta\right])^2\right]\right)^2
\end{aligned}
$$

and simply plugging the sample moments into these expressions. It is well worth noting that the above moment estimators coincide with the maximum likelihood estimators for the OU-Normal model, when one use only the first two moments.

The solution of how to estimate $\mathbb{E}\big[\exp\{\bar{e}_{\Delta t}(\tau)\}\big]$ and $\mathbb{E}\big[\exp\{\bar{e}_{\Delta t}(\tau_1)+\bar{e}_{\Delta t}(\tau_2)\}\big]$ is not as immediate. The residuals $e_{\Delta t}(\tau)$ are effectively random functions and thus complicated to model in general and we therefore suggest to estimate the above expectations non-parametrically. We assume that we are observing m traded maturities $0 < \tau_1 < \tau_2 < \ldots < \tau_m$, also set $\tau_0 = 0$, and that we for each such maturity have observed the forward rate $f_t(\tau_i)$ and the factor process $\mathbf{F}_t$ at time points $t_0 < t_1 < \ldots < t_n$, all a time equal to the investment horizon Δt apart. We may thus calculate the residuals $e_{t_0}(\tau_j), e_{t_1}(\tau_j), \ldots, e_{t_n}(\tau_j)$, for each $j \leq m$. By replacing integrals with sums we approximate the ith integrated residual $\bar{e}_{\Delta t_i}(\tau_j)$ by

$$\begin{aligned}\bar{e}_{\Delta t_i}(\tau_j) \approx & -\sum_{k=1}^{j}(\tau_k-\tau_{k-1})\,\frac{e_{t_i}(\tau_k)-e_{t_{i-1}}(\tau_k)+e_{t_i}(\tau_{k+1})-e_{t_{i-1}}(\tau_{k+1})}{2}\\ & +\Delta t\frac{e_{t_i}(\tau_j)+e_{t_{i-1}}(\tau_j)}{2}.\end{aligned}$$

As an estimate of $\mathbb{E}\big[\exp\{\bar{e}_{\Delta t}(\tau)\}\big]$ we then take the sample mean of the exponent of the above quantities and analogously for $\mathbb{E}\big[\exp\{\bar{e}_{\Delta t}(\tau_1)+\bar{e}_{\Delta t}(\tau_2)\}\big]$.

5.3 Case study

We will in this section estimate the mean-variance optimal portfolio of rolling horizon bonds. To estimate our model we use monthly quotes of euro swap rates from August 2001 to June 2011 provided by Bloomberg Finance LP. We have quotes for maturities $\tau = 1, 2, \ldots, 25$ years. From these we can bootstrap the corresponding zero coupon rates. Since euro swaps have annual payments and we have data for maturities one year apart out to 25 years, we can derive the zero coupon rates without having to resort to any curve fitting or interpolation.

Figure 1 shows the interest rates for some of the maturities in the data. We see that the changes in interest rate level between different maturities are positively correlated and that shorter maturities usually have lower interest rates than longer ones. From these data we can try to obtain the returns from holding the zero coupon bonds for a short period and then rolling them over. Since the data is monthly, we use monthly rolling, and since we do not

Figure 1: Some zero coupon yields derived from euro swap rates. *Source: Bloomberg Finance LP*

have observations of the rates for maturities $\tau - 1/12$ years, we use linear interpolation. Figure 2 shows the cumulative return of all these 25 rolling horizon bonds. At the end the cumulative returns are ordered with higher return for those with longer maturity. The returns for different maturities have positive correlation and longer maturities implies a larger variance of the returns.

It might be instructive to compare the return of a rolling horizon bond with a maturing bond. Figure 3 shows the cumulative return from investing in a bond with 9 years and 10 months left to maturity and either rolling over this bond each month to a new one with 9 years and 10 months to maturity, or keeping it to maturity. At first the returns are similar, but as the remaining time to maturity decreases, the zero coupon bond has progressively lower returns and lower variance of its returns. For good measure the total returns of a European stock index is also shown. At the end of the period the stocks have yielded essentially a zero return and with hindsight bonds were clearly the better investment. Within the sample one also finds that the correlation of the monthly returns from the stocks and the bonds were negative for all maturities, so that bonds would also have provided clear diversification benefits. However, we will in the remainder focus on pure bond portfolios.

In our study we only include rolling horizon bonds. While inclusion of lia-

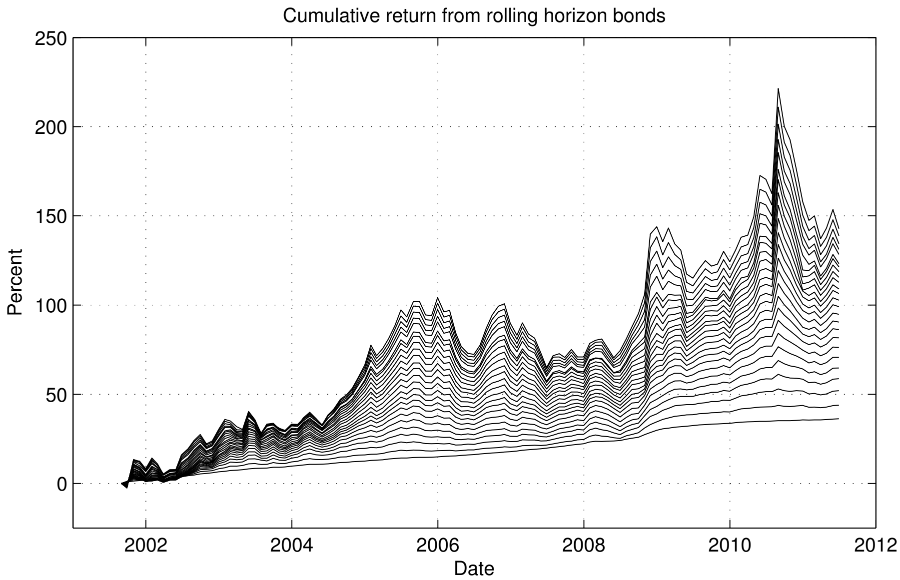

Figure 2: Cumulative return from investing in rolling horizon bonds with time to maturities from 1 to 25 years. *Source: Bloomberg Finance LP*

bilities and other asset classes in the case study would make it more closely comparable to the situation facing an institutional investor. The paper is focused on rolling horizon bonds and we believe that the behavior of the model is more clear if we don't include other assets.

We will compare portfolios based directly on the sample means and covariances of the returns as shown in Figure 2 with portfolios based on a Nelson-Siegel model of the yield curve, where

$$y_t(\tau) = \beta_{0t} + \beta_{1t}\Big(\frac{1-e^{-\gamma\tau}}{\gamma\tau}\Big) + \beta_{2t}\Big(\frac{1-e^{-\gamma\tau}}{\gamma\tau} - e^{-\gamma\tau}\Big).$$

This corresponds to setting

$$\begin{aligned}\kappa_1(\tau) &= 1,\\ \kappa_2(\tau) &= e^{-\gamma\tau},\\ \kappa_3(\tau) &= \gamma\tau e^{-\gamma\tau}.\end{aligned}$$

Here we take $\gamma = 0.7308$, as in [Diebold and Li, 2006] where $\gamma = 0.0609$ with time in months.

For each month we fit the three β parameters using least-squares. The fit is not improved much by chosing another value of γ. The root mean squared

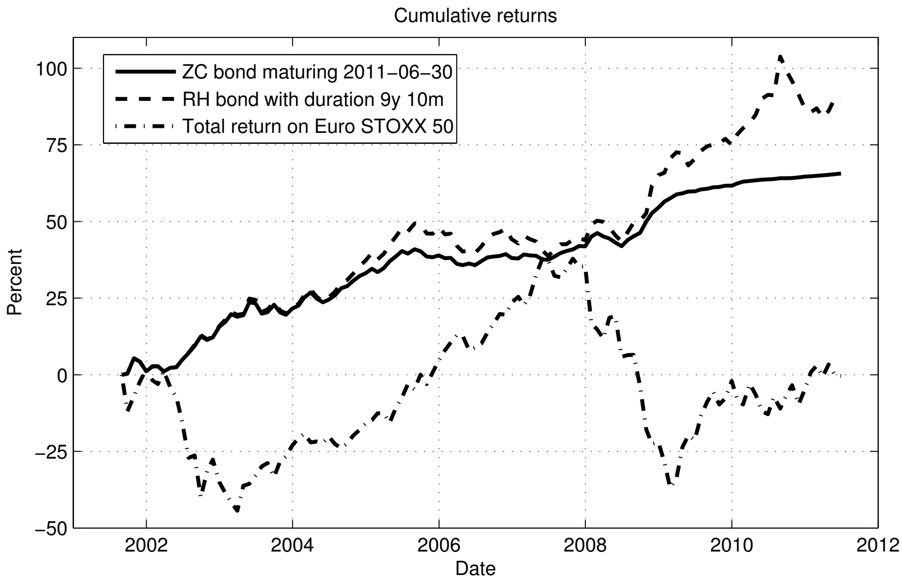

Figure 3: Cumulative returns from investing in a rolling horizon bond, a zero coupon bond, and the stock market. *Source: Bloomberg Finance LP*

errors are of the order a couple of basis points, i.e. hundreds of a percentage point. The time series of the parameters is plotted in Figure 4.

We will assume independence between the time series. The study by [Diebold and Li, 2006] shows that the independence assumption produces better forecasts than a model with dependence, and we will proceed accordingly. Our application and setting is somewhat different from theirs but the findings in the study gives some credence to our assumption. We fit three independent OU-Normal processes to the time series, using the maximum likelihood estimates. The fitted parameters can be found in Table 1.

	β_0	β_1	β_2
$\hat{\mu}(\%)$	4.53	-1.88	-3.27
$\hat{\sigma}^2$ (%)	0.65	1.06	2.74
λ (year^{-1})	0.76	0.35	1.60

Table 1: OU parameter estimates

We will in the case-study not consider the cumulant based estimation approach, as discussed in Section 5.2. This since it turns out that the portfolios obtained from this approach, using as much as the first 6 cumulants, are very

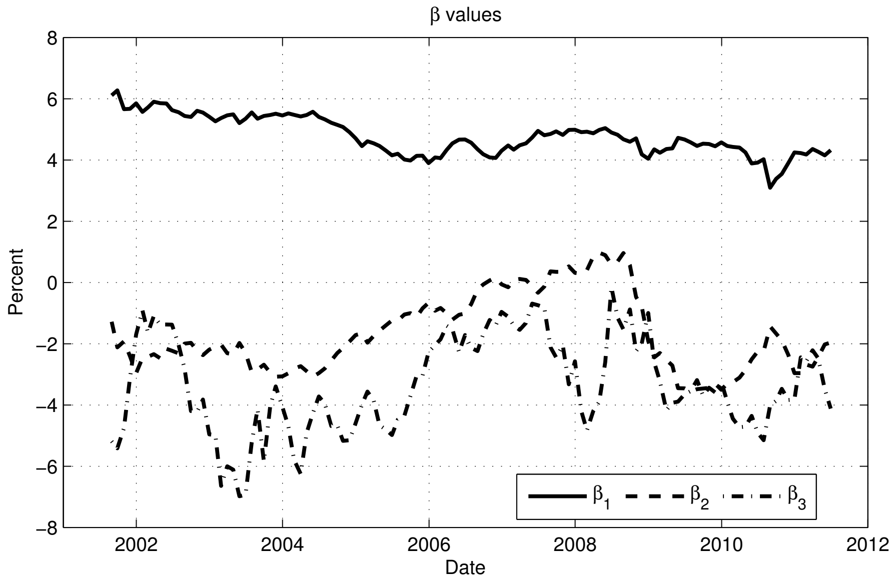

Figure 4: Time series plot of Nelson-Siegel parameters fitted to zero coupon yields.

similar to the ones from the OU-Normal assumption, thus also indicating that this is a fair assumption.

With this, and the estimates of the residuals discussed above, we are able to calculate estimates of the expected return and covariance of rolling horizon bonds based on zero coupon bonds of the above mentioned maturities. For the Nelson-Siegel model, the expected returns of course depends on the initial shape of the yield curve, whereas the expected returns estimated directly from the observed returns do not. Figure 5 shows two yield curves, one so-called "normal" with positive slope which in fact is the mean curve in the sample, and one so-called "inverted" curve where interest rates for shorter maturities are higher than those with longer maturities. Yield curves are seldom inverted, and when they are, it is typically a sign of distress. For example, in our data the yield curve was slightly inverted at the end of 2008 at the height of a financial crisis. Figure 6 shows how the inverted curve, which admittedly is exaggerated in our example, produce different expected returns compared with the normal. Since interest rates for shorter maturities are higher than those with longer, the former are expected to decrease more, and depending on the maturity the expected returns may even be higher for shorter maturities than longer, and longer maturities may have negative expected returns.

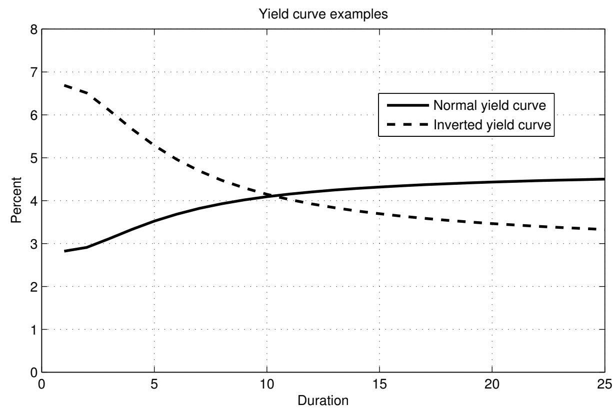

Figure 5: The two initial yield curves considered.

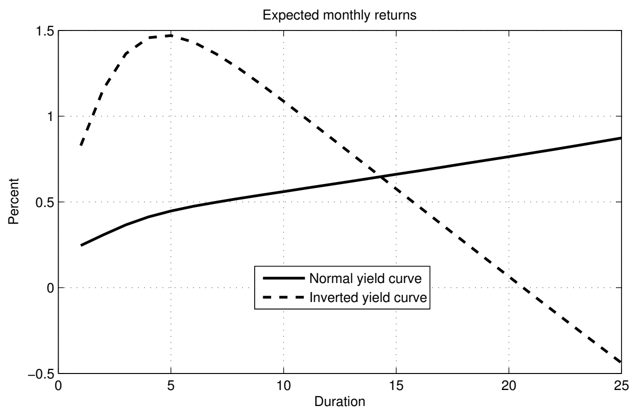

Figure 6: The expected returns of rolling horizon bonds for the two initial yield curves.

The result of the portfolio optimization for a one month horizon, starting from the normal yield curve of Figure 5, and with no short selling, is depicted in Figure 7. We can make some observations from this picture. The middle and lower panel show the composition of each portfolio on the efficient frontier for the method based on the direct returns (let us call it "the empirical portfolio") and the method based on the OU yield curve model ("the Nelson-Siegel portfolio"). In both panels, increasing risk and expected return corresponds to an increase in average time to maturity of the rolling horizon bonds. We also see that the the empirical portfolio seldom has more than three rolling horizon bonds and that the Nelson-Siegel portfolio seldom has more than two. That so few bonds are included in a portfolio with given risk is reasonable since 99.7% of all variation of the returns is explained by the three first factors in a principal component analysis. That fewer are needed in the Nelson-Siegel model may be due to the fact that it imposes more rigidity to the possible changes in the yield curve.

In the case of the inverted yield curve, the Nelson-Siegel model produces portfolios according to the same pattern as for the normal yield curve, but with the maximum maturity being about 5 years since the expected return decreases when increasing the maturity further.

For a one year horizon the composition of the empirical portfolios along the efficient frontier is essentially the same as the one for the one month horizon, as seen in Figure 8. The main difference is that the expected returns and volatility has increased due to the longer time horizon. For the Nelson-Siegel portfolios there are some changes in composition for lower returns and volatilities, but it is still the case that few bonds are needed in order to achieve an efficient portfolio. There is however a difference when one compares the efficient frontiers of both models. The Nelson-Siegel frontier does not reach as far as the empirical one since the assumption of mean reversion decreases the expected return of bonds with longer maturity. The mean reversion also decreases the volatility of the returns.

Figure 9 shows the frontier for the Nelson-Siegel portfolios when short selling is allowed. In particular, the portfolios have been constrained to have zero net position with maximum 50% long and 50% short position. Portfolios with higher risk and expected return are funded by being short the 1 year bond and long bonds with longer maturities. (This position is usually called a flattener since it makes better returns as the yield curve flattens. For less risky portfolios the model prescribes portfolios that are long bonds with intermediate maturities and are short bonds with both shorter and longer maturities. This portfolio makes better returns if the curvature of the yield curve decreases.)

In-sample returns can be deceiving so we also want to check the out-of-sample performance. Since the Nelson-Siegel model takes the current shape of the yield curve into account — in contrast to the empirical portfolios — one may hope that it would perform better over time. We estimate both models for every month with a three year rolling window of data and check the performance of the portfolio with the maximum expected return, the minimum expected return and 8 portfolios with equally spaced intermediate expected returns. We produce the efficient frontiers for both a one month horizon and a one year horizon. We also include the performance of the portfolio that ex ante would perform the best if the yield curve is unchanged from one month to the next. In fixed income parlance, it maximizes the carry

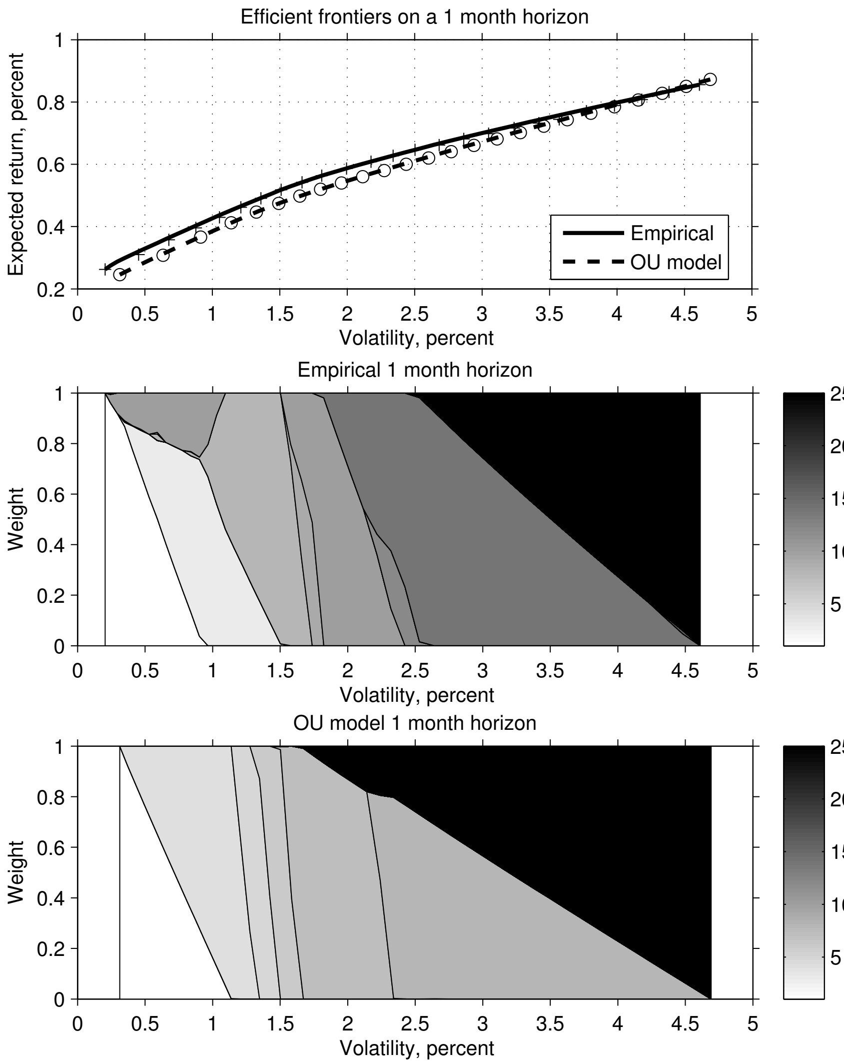

Figure 7: Top figure shows efficient frontiers for a one month horizon together with the empirical (crosses) and OU model (circles) individual rolling horizon bonds included in the optimization. Bottom two figures show corresponding portfolios.

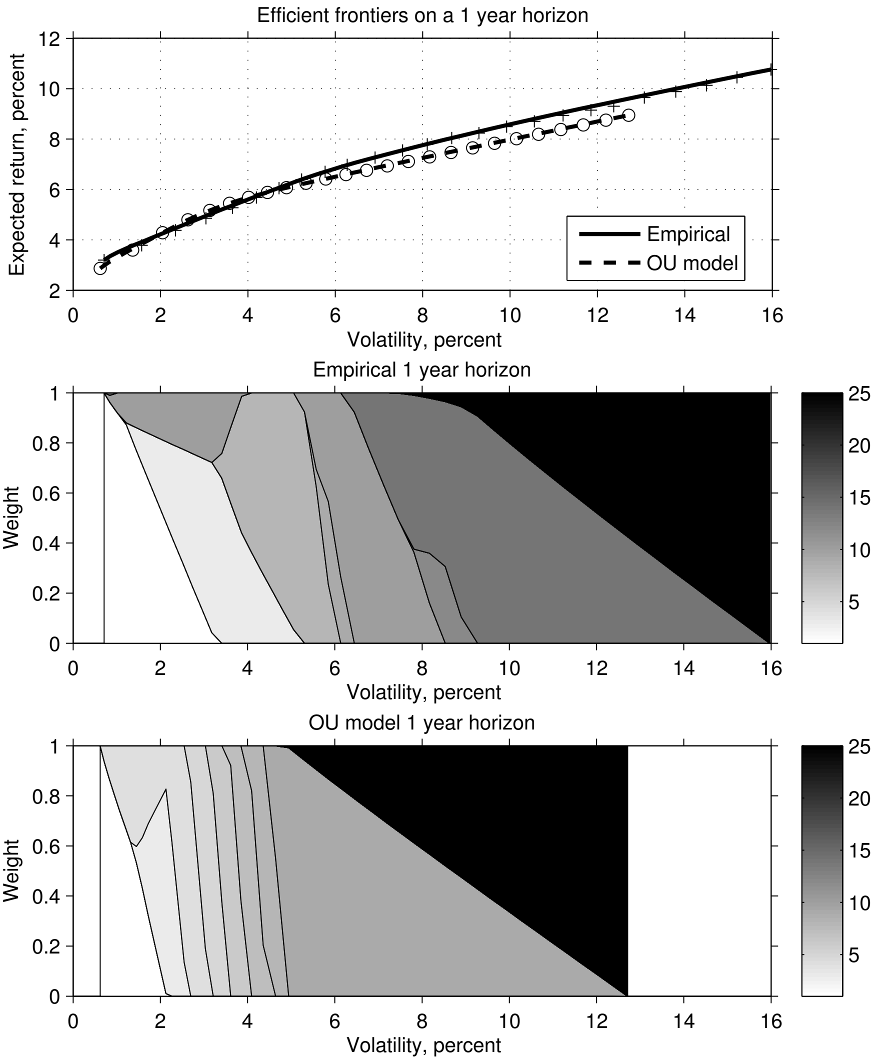

Figure 8: Top figure shows efficient frontiers for a one year horizon together with the empirical (crosses) and OU model (circles) individual rolling horizon bonds included in the optimization. Bottom two figures show corresponding portfolios.

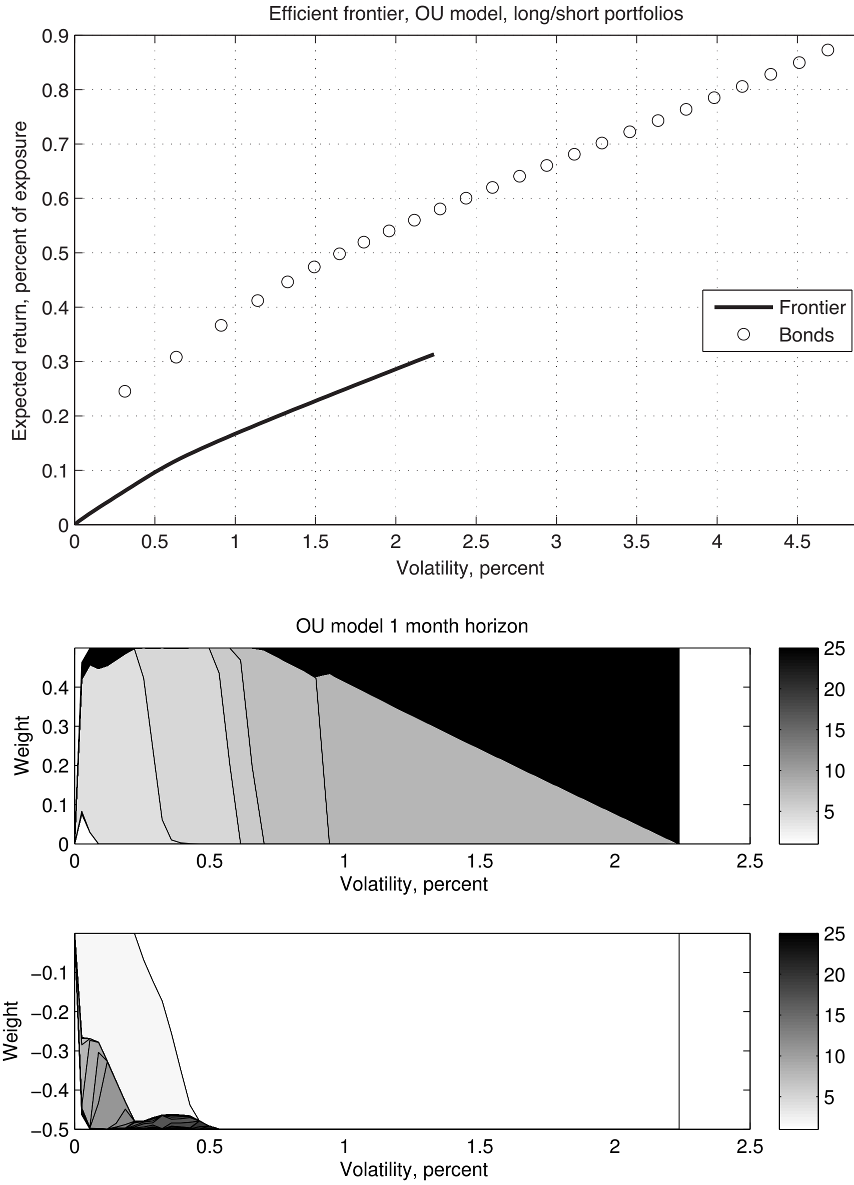

Figure 9: Top figure shows efficient frontiers for a one month horizon with zero net position together with the OU model individual rolling horizon bonds (circles) included in the optimization. Bottom two figures show corresponding portfolios.

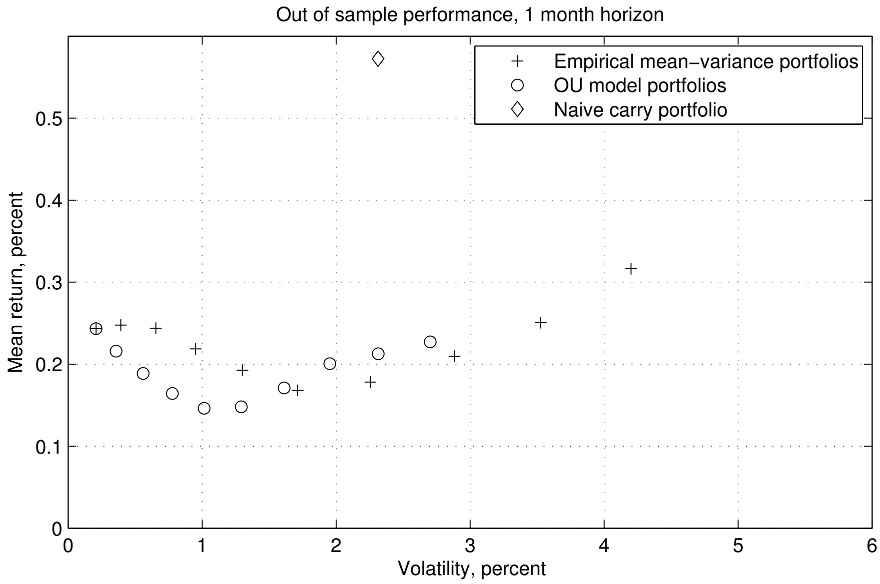

Figure 10: Out-of-sample performance on a one month horizon. The ten plusses and circles correspond to the ten portfolios on the efficient frontier for the empirical and Nelson-Siegel model, respectively.

(incl. rolldown).

Figure 10 shows the abysmal performance of both the empirical portfolios and the Nelson-Siegel methods for the one month horizon. The relation between mean return and volatility is not even increasing. The naive carry portfolio which does not take any history into account other than the starting curve handily outperforms both of the other two methods. Figure 11 shows the cumulative return of the maximum expected return portfolios for the different strategies. It is clear that the volatility of both the empirical and the Nelson-Siegel portfolios is high. It is disappointing that the Nelson-Siegel model is unable to capitalize on its knowledge of the starting yield curve. One possible conclusion is that the mean reversion is too low to be of use over a one month horizon.

On a one year horizon things look better as shown in Figure 12, but the naive carry method still outperforms the two other methods. The returns that underlies Figure 12 come from overlapping series since the portfolios are reestimated each month.

One can see the effect of different starting points in Figure 13. It shows the twelve different performances one get when starting estimation in each of the twelve months of the year. One of the marks represents estimating

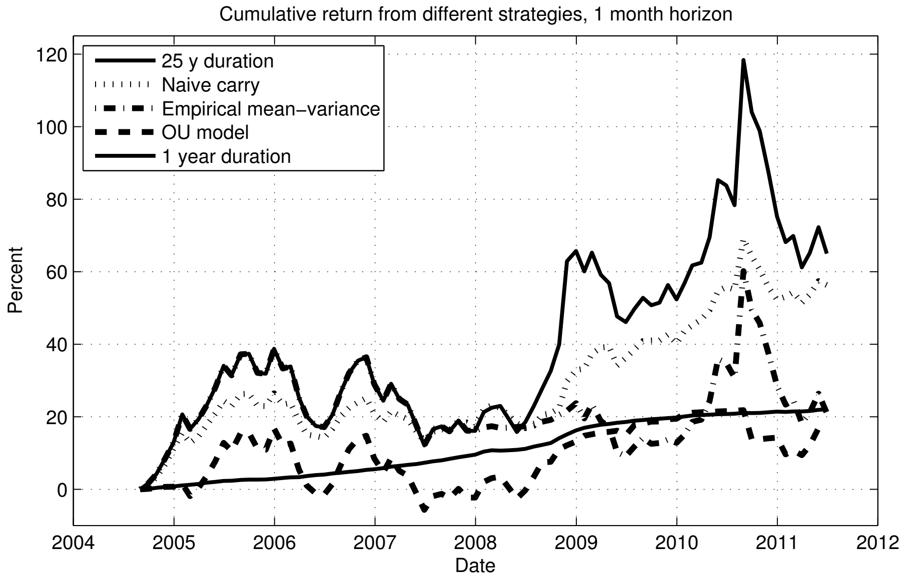

Figure 11: Cumulative returns from the maximum return portfolios from the empirical and Nelson-Siegel model, together with the maximum carry portfolio and the rolling horizon bonds with 1 year and 25 years to maturity.

both portfolios with data from January to December and then letting both run for a year and then repeating the estimation with the last year's data. Another mark represents estimating the portfolios with data from February to January and then lettig both run for twelve months, and then repeating the estmation with the new data. The remaining ten marks represent the respective results when starting from March, . . . , December, respectively. We see that the Nelson-Siegel method has lower variance of its returns compared to the empirical mean-variance method.

That the mean-variance portfolio of rolling horizon bonds performs so poorly out of sample should be understood as a failure of the mean-variance approach and not as a failure of the rolling horizon bonds. This pattern is well known from e.g. equites. In a real application one would have to use more advanced methods, see for example [Black and Litterman, 1992].

6 Conclusion

We have studied so called rolling horizon bonds, where the time to maturity is held fixed, in particular under the assumption of an affine term structure

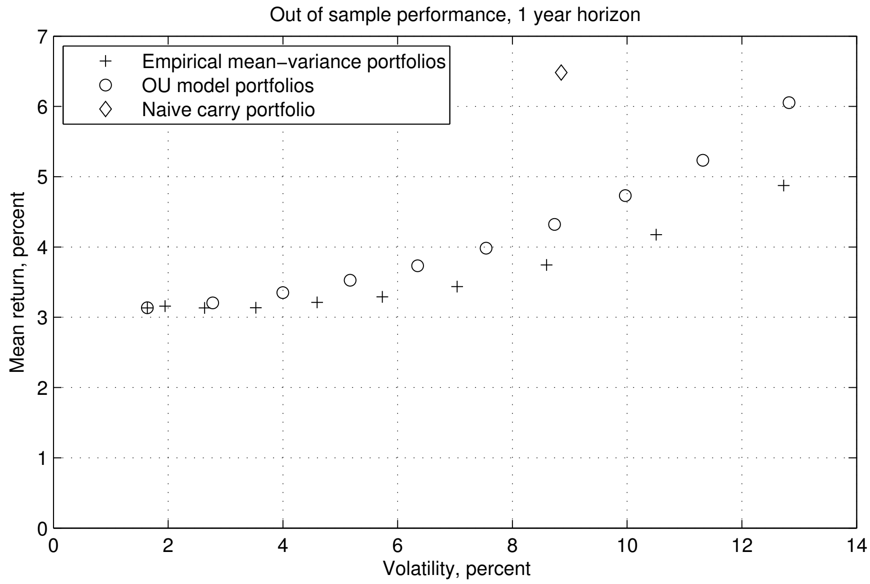

Figure 12: Out-of-sample performace on a one year horizon. The ten plusses and circles correspond to the ten portfolios on the efficient frontier for the empirical and Nelson-Siegel model, respectively.

where the stochastic processes are generalised OU processes. We show that the return and variance of a portfolio of such bonds can be expressed using the cumulant generating functions associated with the generalised OU processes. This allows for calculation of the efficient mean-variance portfolios and we also provide some suggestions for the estimation of the unknown parameters. The results can be used by an investor seeking to invest in bonds with a fixed time to maturity, for example a pension fund.

In a case study using euro swap rates we see that the estimated mean-variance portfolios does not perform well in out of sample testing. This should however be attributed to the failures of the mean-variance procedure when the market parameters need to estimated. In a real application more advanced portfolio selection methods should be used. But since the mean-variance approach is a standard method we believe it serves to illustrate the usefulness of the first part of the paper.

We have also seen that if we assume that the short-rate is given by an affine relation in a risk-neutral setting the term-structure is given by the cumulant generating functions associated with the generalised OU process. This provides an intuitive way of examining the possible term structures that can be generated by a certain short rate dynamics.

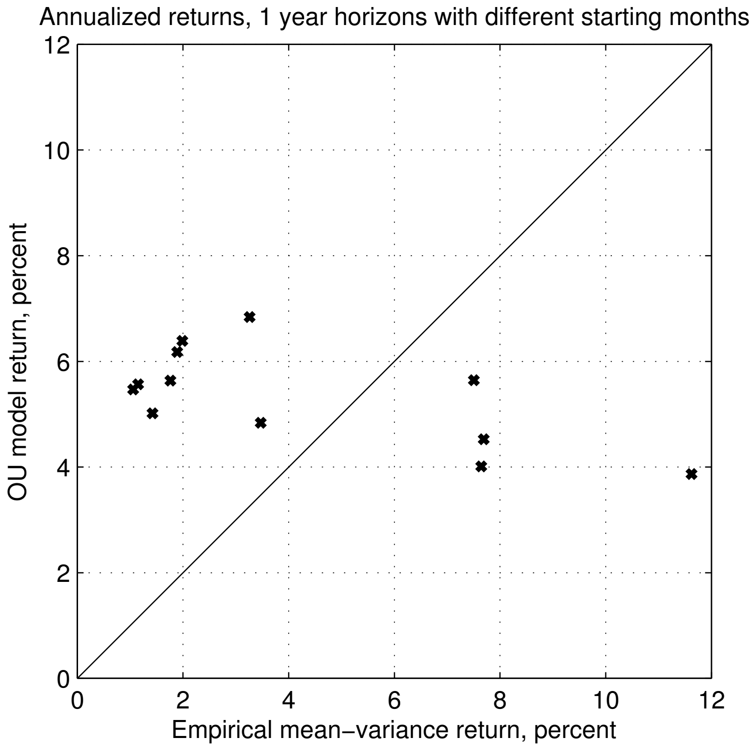

Figure 13: Cumulative annualized returns for the maximum return portfolios for the empirical model against the Nelson-Siegel model for 12 different starting months. The Nelson-Siegel model gives the better return for 8 out of 12 months, and the range of returns is markedly lower than those of the empirical model.

Acknowledgements

Patrik Andersson's research was in part conducted under the fellowship program - JSPS Postdoctoral Fellowship for North American and European Researchers (Short Term). Andreas Lagerås's research is financially supported by AFA Insurance; the views and opinions are however not necessarily those of AFA Insurance. The authors are grateful to Bloomberg Finance LP for providing data and to the anonymous referees.

References

[Barndorff-Nielsen and Shephard, 2003] Barndorff-Nielsen, O. E. and Shephard, N. (2003). Integrated OU processes and non-Gaussian OU-based stochastic volatility models. *Scand. J. Statist.*, 30(2):277–295.

[Black and Litterman, 1992] Black, F. and Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43.

[Coroneo et al., 2008] Coroneo, L., Nyholm, K., and Vidova-Koleva, R. (2008). How arbitrage-free is the Nelson-Siegel term model? *ECB Working paper*, 874.

[Diebold and Li, 2006] Diebold, F. X. and Li, C. (2006). Forecasting the term structure of government bond yields. *J. Econometrics*, 130(2):337–364.

[Duffie et al., 2003] Duffie, D., Filipović, D., and Schachermayer, W. (2003). Affine processes and applications in finance. *Ann. Appl. Probab.*, 13(3):984–1053.

[Ekeland and Taflin, 2005] Ekeland, I. and Taflin, E. (2005). A theory of bond portfolios. *Ann. Appl. Probab.*, 15(2):1260–1305.

[Filipović, 1999] Filipović, D. (1999). A note on the nelson-siegel family. *Math. Finance*, 9(4):349–359.

[Piazzesi, 2009] Piazzesi, M. (2009). Affine term structure models. In Aït-Sahalia, Y. and Hansen, L., editors, *Handbook of Financial Econometrics, Vol. 1.* North Holland.

[Rutkowski, 1999] Rutkowski, M. (1999). Self-financing trading strategies for sliding, rolling-horizon, and consol bonds. *Math. Finance*, 9(4):361–385.