

WP 2004-8

Dynamic asset allocation and latent variables

by

Carsten Sørensen and Anders Bjerre Trolle

**INSTITUT FOR FINANSIERING, Handelshøjskolen i København
Solbjerg Plads 3, 2000 Frederiksberg C
tlf.: 38 15 36 15 fax: 38 15 36 00**

**DEPARTMENT OF FINANCE, Copenhagen Business School
Solbjerg Plads 3, DK - 2000 Frederiksberg C, Denmark
Phone (+45)38153615, Fax (+45)38153600
www.cbs.dk/departments/finance**

ISBN 87-90705-87-4
ISSN 0903-0352

Dynamic asset allocation and latent variables

Carsten Sørensen and Anders Bjerre Trolle

Department of Finance
Copenhagen Business School
Solbjerg Plads 3
DK-2000 Frederiksberg
Denmark

JEL Classification: G11

Keywords: Portfolio choice, predictability, VAR, unobserved state-variables, hedging demands

December 2004

Acknowledgements: Carsten Sørensen gratefully acknowledges financial support from the Danish Social Science Council.

Dynamic asset allocation and latent variables

Abstract

We derive an explicit solution to the portfolio problem of a power utility investor with preferences for wealth at a finite investment horizon. The investor can invest in assets with return dynamics described as part of a general multivariate model. The modeling framework encompasses discrete-time VAR-models where some of the state-variables (e.g. expected excess returns) may not be directly observable. A realistic multivariate model is estimated and applied to analyze the portfolio implications of investment horizon and return predictability when real interest rates and expected excess returns on stock and bonds are not directly observed but must be estimated as part of the problem faced by the investor. The solution exhibits small variability in portfolio allocations over time compared to the case when excess returns are assumed observable.

1 Introduction

The solution to a multi-period portfolio problem can differ substantially from the solution to a static or single-period portfolio problem, as demonstrated originally by Samuelson (1969) and Merton (1969,1971,1973). This paper offers an explicit solution to a basic multi-period dynamic portfolio problem when the return dynamics are described by a multivariate time-series model and the investor is concerned with maximizing the expected utility of wealth at a given horizon. The modeling framework encompasses return generating models where some of the basic state-variables are unobserved, and where the investor is faced with a filtering problem as part of the overall dynamic asset allocation problem. Our solution makes it possible to address, e.g., the portfolio implications of an estimated VAR-model that involves return-predictability for investors with different risk aversion and time horizons in a simple and consistent manner. Furthermore, when some of the state-variables are unobserved, we establish a close link between how unobserved state-variables can be handled consistently in the econometric estimation of the model as well as in a subsequent analysis of optimal asset allocation choice by using a Kalman filtering approach. This is explored in a realistic model calibration.

The multivariate discrete-time modeling of return dynamics is basically similar to the multivariate VAR-setting used by Campbell, Chan and Viceira (2003), but extended to the situation where some of the state-variables may not be directly observed by the investor. The general version of our return generating model is based on a state-space representation which consists of a *transition* equation and a *measurement* equation. The transition equation describes the return dynamics, and this is exactly the Campbell et al. (2003) multivariate VAR-model. The measurement equation describes what is being observed (and what is not).

The Campbell et al. (2003) return generating model is the special case where all state-variables are directly observed. Campbell et al. (2003) investigates the optimal asset allocation and consumption policy of an infinitely-lived investor with Epstein-Zin recursive preferences, and relies on an approximate solution methodology in order to solve for the optimal policies. Also, Campbell et al. (2003) must assume that the VAR-model is time-homogenous in order to solve the infinite-horizon investment and consumption problem. Due to the more simple assumption of power-utility of terminal wealth, we can relax these assumptions and still address the effects of time horizon and risk aversion on optimal asset allocation. When state-variables are directly observed, we thus obtain an explicit solution to the relevant dynamic portfolio problem and a simple recursive solution algorithm for implementing the solution. The involved

recursive solution algorithm solves a particular system of difference equations; this system of difference equations is analogous to the multi-dimensional Riccati equation that arises in related continuous-time contexts when solving the relevant Hamilton-Jacobian-Bellman equation; see e.g. Liu (1999).

When some of the state-variables are not directly observed, the same solution procedure applies. However, in this case the relevant state-variables are now the perceived values of the possibly latent state-variables. The perceived values of the state-variables can recursively be found by Kalman filtering, which is applicable to any model that has a state-space representation. As we demonstrate, the dynamics of the Kalman filtered state-variables is also described by a VAR-model, and the solution methodology and recursive algorithm used for the case without observation noise can also be applied to this general case.

In a continuous-time framework, Williams (1977), Detemple (1986), Dothan and Feldman (1986), and Gennotte (1986) have previously demonstrated that the dynamic portfolio problem of an investor, who cannot directly observe the state-variables, separates into a filtering problem, in which the investor estimates the state-variable position, and an investment problem where the filtered estimates are treated as the relevant state-variables. Furthermore, in a setting with a single risky asset, Detemple (1986) and Gennotte (1986) show that the uncertainty about the instantaneous excess return on the risky asset is not affecting the optimal portfolio choice, which simply takes the form of Merton (1969) by substituting the instantaneous expected excess return with its perceived value. While our discrete-time solution also allows the filtering problem and the investment problem to be handled separately and consecutively, the uncertainty about the exact positions of the unobserved state-variables affects the optimal portfolio choice in our discrete-time setting. For example, even in a one-period model with a single risky asset, the relevant variance that must be used along with the perceived value of the expected excess return in order to determine the optimal portfolio (in the Markowitz (1952), Merton (1969), and/or Samuelson (1969) formulas) is affected by the uncertainty about the true expected excess return. In such a one-period setting, our solution resembles and coincides with the Bayesian approach of incorporating parameter uncertainty into the portfolio choice problem, as originally carried out by e.g. Klein and Bawa (1976) and Bawa, Brown and Klein (1979).

In a dynamic context, the discrete-time solution with uncertainty about time-varying expected excess returns in general gives rise to less risky investments and less variability in portfolio allocations over time compared to the case without uncertainty. The tendency for a

decrease in risky investments follows in part from the previous discussion of the similarity with a Bayesian portfolio approach, although this effect seems small for short rebalancing intervals. The decrease in variability in portfolio allocations over time is due to the phenomenon that the process of the filtered state-variables is usually less variable over time than the true process. From Merton (1971,1973) it is well-known that the optimal portfolio allocation of a multi-period investor can be decomposed into a myopic single-period component and a term that describes how the investor should hedge changes in the basic state-variables that describe the opportunity set. Due to the decrease in the variability of the relevant perceived state-variable process, the hedging demand will generally tend to be smaller with uncertainty about latent state-variables than without.

Besides calibrating a realistic multivariate model, we also consider two illustrative and relatively simple examples within the general modeling framework. The two examples are in part included to support the above mentioned qualitative statements about the implications for optimal portfolio allocations and furthermore to illustrate modeling possibilities and features of the discrete-time solution compared to e.g. similar continuous-time solutions under given and specific premises. The first illustrative example considers the effect of learning about the expected return on a stock. The model is similar to a model analyzed numerically by Barberis (2000), and it is a discrete-time version of the model considered by Brennan (1998) in continuous-time. The second illustrative example considers a model where the expected excess return on the stock is assumed mean-reverting. The expected excess return is assumed unobservable and must be filtered based on observed stock returns. The model is a discrete-time version of a continuous-time model considered by Rodriguez (2002) and the basic modeling builds on Kim and Omberg (1996).

This paper introduces latent state-variables into a discrete-time set-up which is otherwise conceptually very similar to the continuous-time settings analyzed by Liu (1999), and many others.¹ It should be noted that since data is usually observed discretely, the discrete-time formulation is directly suitable for econometric purposes. It is, however, generally possible to translate a multivariate continuous-time VAR(1) model to a discrete-time VAR(1) model,² and

¹For example, Kim and Omberg (1996) and Wachter (2002) analyze the effects of time-varying expected excess returns, and Brennan and Xia (2000,2002) and Sørensen (1999) analyze the effects of stochastic interest rates or inflation. Campbell and Viceira (1999,2001,2002) introduce similar features in discrete-time asset allocation models.

²Campbell, Chacko, Rodriguez and Viceira (2004) provide an example. Note, however, that it is generally

this feature is applied in our calibration analysis in order to consistently model the relation between interest rate dynamics and the forms of the term structure of interest rates using a standard continuous-time no-arbitrage approach. In addition, the calibrated model features stochastic inflation and stochastic risk premia on stock and bonds. The model is calibrated to US bond, stock and inflation data for the period 1951 to 2004. The calibrated model is related to models calibrated in e.g. Campbell et al. (2003), Campbell and Viceira (1999,2001), Brennan and Xia (2002,2004), and Sangvinatsos and Wachter (2004), but with the added feature of real interest rates and expected excess returns being treated as unobservable in both the econometric estimation approach and the subsequent analysis of implications for optimal asset allocation.

More specifically, following Sangvinatsos and Wachter (2004), the modeling of interest rate dynamics is based on a specific two-factor essentially affine term structure model, where bond yields and expected excess bond returns are affine in the two term structure state-variables. Expected excess returns on stock are also assumed affine in the term structure state-variables as well as an extra stock specific state-variable. The three state-variables are not observable directly, but indirectly through observations of bond yields and the dividend-price ratio. The theoretical relation between the dividend-price ratio and the state-variables is based on the Campbell and Shiller (1988a,1988b) approximation, which links the current dividend-price ratio to expected future dividend growth and expected future stock returns. Assuming constant expected dividend growth, we thus show that the dividend-price ratio is affine in the three state-variables. The use of a dynamic dividend discount model for estimating expected excess stock returns is inspired by Brennan and Xia (2004), although our implementation is somewhat different.

The specific model is represented in state-space form and estimated by a standard panel data approach. The measurement equation which relates the basic state-variables to bond yields and the dividend-price ratio incorporates “noise” in all relations. Therefore, the latent state-variables are never perfectly observable. Having calibrated parameters in the model, the not possible to translate a discrete-time VAR(1) to a continuous-time VAR(1) model without loss of generality. Consider for example the discrete-time univariate AR(1) model : $x_{t+1} = \alpha x_t + \epsilon_{t+1}$. The discrete-time solution to the continuous-time counterpart, the Ornstein-Uhlenbeck process, is also a VAR(1) model, but with the restriction $\alpha > 0$. This parameter restriction can be a limitation since the AR(1) process with $\alpha \leq 0$ may be a reasonable specification in some contexts since it includes the “white noise” process ($\alpha = 0$) and other reasonable stationary processes ($-1 < \alpha < 0$).

model with observation “noise” is also used to determine optimal policies for an asset allocator. Most importantly, the calibrations show that optimal stock allocations vary significantly less over time compared to similar studies that assume that e.g. expected stock returns are directly observable.

The paper is organized as follows: We set up the basic multivariate model in Section 2 where all state-variables are assumed observed. In Section 3 the model and solution are extended to the case where some state-variables are not directly observed. In order to analyze and pinpoint how uncertainty about unobserved state-variables affect the optimal portfolio allocation, two illustrative examples are presented in Section 4. Finally, the model used for calibration and the empirical analysis is presented in Section 5, while Section 6 concludes. An Appendix contains details and proofs.

2 The basic multivariate model of portfolio choice

In this section we consider a special case of the general state-space model. The discrete-time return dynamics are specialized to a VAR(1) model, as was also considered by Campbell et al. (2003) in a similar portfolio context. In many respects, the portfolio solution also shares similar features with the approximate solution obtained by Campbell et al. (2003) for an infinitely long-lived investor with Epstein-Zin recursive preferences. In order to facilitate comparison with the results in Campbell et al. (2003), the notation and model of asset return dynamics in this section is therefore basically adopted from their setting, and then subsequently extended in the following section.

2.1 Preferences

As in Kim and Omberg (1996), Brennan, Schwartz and Lagnado (1997), Sørensen (1999), Barberis (2000), among others, the investor is endowed with initial wealth, W_0 , which is to be invested to maximize expected power utility of the form

$$E_0[U(W_T)] \quad \text{where} \quad U(W) = \frac{W^{1-\gamma} - 1}{1-\gamma} \quad (1)$$

and where $\gamma > 0$ is the parameter of constant relative risk aversion, and T is the investment horizon. For $\gamma = 1$, we have the logarithmic utility function, $U(W) = \log W$, as a limiting special case. The investor is only concerned with maximizing the utility of terminal wealth and is assumed not to use wealth for intermediate consumption, nor accumulate additional

wealth due to labor income. Basically, the asset allocation problem may be thought of as the problem faced by an individual who has received a lump sum that must be invested for the purpose of retirement at time T .

2.2 Dynamics of investment returns

The investor can invest in n assets, and the dynamics of the relevant asset returns and state-variables are described by a first-order vector autoregressive process, VAR(1). Following Campbell et al. (2003), let $R_{i,t+1}$ denote the real gross return on asset i , and let $r_{i,t+1} = \log(R_{i,t+1})$ denote the similar real log return on asset i ($i = 1, \dots, n$). The relevant state-vector in the analysis is then given in the following stacked form:

$$z_{t+1} = \begin{pmatrix} r_{1,t+1} \\ x_{t+1} \\ s_{t+1} \end{pmatrix} \quad (2)$$

where $x_{t+1} = ((r_{2,t+1} - r_{1,t+1}), (r_{3,t+1} - r_{1,t+1}), \dots, (r_{n,t+1} - r_{1,t+1}))'$ is an $(n - 1) \times 1$ vector of log excess returns, and s_{t+1} is a vector of other relevant state-variables (including e.g. the dividend-price ratio). The state vector, z_{t+1} , has dimension $m \times 1$. The excess returns are measured relatively to asset 1 which may be thought of as a short risk-free asset. However, the realized return, or ex post return, on asset 1 is more generally allowed to be stochastic. This is especially relevant when the investor does not have access to a risk-free asset in real terms, and $r_{1,t+1}$ instead refers to the realized real return on, say, a short nominal treasury bill, as in our subsequent empirical application of the portfolio choice model.

The VAR(1) model of state-variable dynamics is given by

$$z_{t+1} = \Phi_0 + \Phi_1 z_t + v_{t+1} \quad (3)$$

where Φ_0 is an m -dimensional vector and Φ_1 is an $m \times m$ matrix. The innovations v_{t+1} are assumed to be uncorrelated and identically normally distributed:

$$v_{t+1} \sim \mathcal{N}(0, \Sigma),$$

where

$$\Sigma = \text{Var}_t(v_{t+1}) = \begin{pmatrix} \sigma_1 & \sigma'_{1x} & \sigma'_{1s} \\ \sigma_{1x} & \Sigma_{xx} & \Sigma'_{xs} \\ \sigma_{1s} & \Sigma_{xs} & \Sigma_{ss} \end{pmatrix}. \quad (4)$$

2.3 Portfolio returns

The solution is based on the following characterization of the log return on the investment portfolio:

$$r_{p,t+1} = r_{1,t+1} + \alpha_t' x_{t+1} + \frac{1}{2} \alpha_t' (\sigma_x^2 - \Sigma_{xx} \alpha_t), \quad (5)$$

where α_t is an $(n-1)$ -dimensional vector of portfolio weights that the investor holds in asset 2 to asset n in the period between t and $t+1$, and where $\sigma_x^2 = \text{diag}(\Sigma_{xx})$ is the vector of diagonal elements of the excess return variance-covariance matrix, Σ_{xx} . The expression for the portfolio log return in (5) can be shown to be exact for small time intervals (i.e. in continuous-time), or whenever one interprets α as the constant portfolio weights maintained in the interval between t and $t+1$ by continuously adjusting the portfolio in the interval; however, if the investor is assumed to follow a passive buy-and-hold strategy between t and $t+1$, then (5) is only an approximation of the log portfolio return over the interval. The approximation is also used and discussed in Campbell and Viceira (2002) and Campbell et al. (2003).

2.4 Solving the model

Let V_t denote the value function of the utility maximization problem faced by the investor. Then the value function, V_t , and the optimal portfolio policy, α_t , must at any discrete time point t satisfy the Bellman equation,

$$V_t = \max_{\alpha_t} E_t [V_{t+1}], \quad t = 0, \dots, T-1, \quad (6)$$

and with $V_T = U(W_T)$. In order to solve the model for the optimal portfolio policy and the form of the value function, we make the following conjecture about the functional forms these take:

$$\alpha_t = A_0(t) + A_1(t) z_t \quad (7)$$

where $A_0(t)$ is an $(n-1)$ -dimensional vector and $A_1(t)$ is an $(n-1) \times m$ matrix. Furthermore,

$$V_t = \frac{\left(e^{B_0(t) + (1-\gamma)B_1(t)'z_t + (1-\gamma)z_t' B_2(t)z_t} \right) W_t^{1-\gamma} - 1}{1-\gamma} \quad (8)$$

where $B_0(t)$ is a scalar, $B_1(t)$ is an m -dimensional vector, and $B_2(t)$ is a symmetric $m \times m$ matrix. In particular, at the terminal date we must have that $B_0(T) = 0$, $B_1(T) = 0$, and $B_2(T) = 0$ in order to ensure that $V_T = U(W_T)$.

To verify the above conjecture,³ one can substitute the expression for V_{t+1} into the Bellman equation (6), and make use of the fact that wealth evolves according to $W_{t+1} = W_t e^{r_{p,t+1}}$, where the log portfolio return, $r_{p,t+1}$ is given in (5). Then, by evaluating the expectations involved in the Bellman equation and maximizing the resulting expression with respect to α_t , one obtains optimal portfolio weights that takes the form in (7), where

$$A_0(t) = \Omega(t+1)^{-1} \left[H_x \Gamma(t+1)' \Phi_0 + \frac{1}{2} \sigma_x^2 + (1-\gamma) H_x \Sigma \Gamma(t+1) (B_1(t+1) + H_1') \right], \quad (9)$$

$$A_1(t) = \Omega(t+1)^{-1} H_x \Gamma(t+1)' \Phi_1. \quad (10)$$

and where H_1 and H_x are selection matrices that select respectively the first element and a vector of the next $n-1$ elements from an m -dimensional vector. (In particular, $r_{1,t+1} = H_1 z_{t+1}$ and $x_{t+1} = H_x z_{t+1}$.) Furthermore, in the expressions for $A_0(t)$ and $A_1(t)$ we have made use of the auxiliary matrices,

$$\Omega(t+1) = \Sigma_{xx} + (\gamma-1) H_x \Sigma \Gamma(t+1) H_x', \quad (11)$$

$$\Gamma(t+1) = (I_m + 2(\gamma-1) B_2(t+1) \Sigma)^{-1}, \quad (12)$$

where I_m denotes the $m \times m$ dimensional identity matrix.

Now, substituting the optimal portfolio described by (7), (9), and (10) back into the right hand side of the Bellman equation (6) and evaluating the expectations, one obtains V_t in the form conjectured in (8), where

$$\begin{aligned} B_0(t) = & B_0(t+1) + \frac{1}{2} \ln |\Gamma(t+1)| + \frac{1}{2} (1-\gamma) A_0(t)' (\sigma_x^2 - \Sigma_{xx} A_0(t)) \\ & + \frac{1}{2} (1-\gamma)^2 (B_1(t+1)' + H_1 + A_0(t)' H_x) \Sigma \Gamma(t+1) (B_1(t+1) + H_1' + H_x' A_0(t)) \\ & + (1-\gamma) \Phi_0' \Gamma(t+1) (B_1(t+1) + H_1' + H_x' A_0(t) + B_2(t+1) \Phi_0), \end{aligned} \quad (13)$$

$$\begin{aligned} B_1(t) = & A_1(t)' \left(\frac{1}{2} \sigma_x^2 + H_x \Gamma(t+1)' \Phi_0 + (1-\gamma) H_x \Sigma \Gamma(t+1) (B_1(t+1) + H_1') \right) \\ & + \Phi_1' \Gamma(t+1) (B_1(t+1) + H_1' + 2B_2(t+1) \Phi_0), \end{aligned} \quad (14)$$

$$B_2(t) = \frac{1}{2} A_1(t)' \Omega(t+1) A_1(t) + \Phi_1' \Gamma(t+1) B_2(t+1) \Phi_1. \quad (15)$$

This ends the verification proof in the sense that we have proven that the portfolio policy in (7) and the value function in (8) constitutes a solution to the Bellman equation (6) when the

³Details of the described verification procedure are given in Appendix A.

time-dependent matrices $A_0(t)$, $A_1(t)$, $B_0(t)$, $B_1(t)$, and $B_2(t)$ satisfy the difference equation system (9)–(15). Moreover, the relevant time-dependent matrices involved in the solution are easily obtained recursively by starting with: (i) setting $B_0(T) = 0$, $B_1(T) = 0$, and $B_2(T) = 0$, (ii) then insert these in (9) and (10) (and (11) and (12)) to obtain the optimal portfolio policy at time $T - 1$, (iii) then insert the expressions for $A_0(T - 1)$ and $A_1(T - 1)$ in (13), (14), and (15) in order to obtain the value function at time $T - 1$. The procedure can now be repeated recursively to determine the solution at time $T - 2$, $T - 3$, and so on until time $t = 0$.

It may be noted that in order to ensure that the auxiliary matrices Ω and Γ defined in (11) and (12) are indeed invertible, as implicitly assumed in constructing the recursive difference equation system (9)–(15), it suffices to ensure that $B_2(t)$ is positive semi-definite (for all $t = 0, \dots, T$). Furthermore, since the variance-covariance matrix Σ_{xx} is always positive definite, and $B_2(T) = 0$ is positive semi-definite, it can be inferred by recursive inspection of (15), (11), and (12) that $\gamma \geq 1$ is a sufficient (but not necessary) condition to ensure this.⁴ For any return parameters, the solution therefore applies to a set of investors that includes the log-utility investor and any more risk averse investor. This is the interesting set of investors since it is well-known that log-utility investors do not hedge changes in the investment opportunity set, while investors that are more risk averse will exhibit portfolio strategies that reflect a desire to hedge undesirable changes in investment opportunities, as originally described by Merton (1971,1973) in a continuous-time setting.

3 The general multivariate model

The general version of the return generating model can be written in state-space form as

$$z_{t+1} = \Phi_0(t) + \Phi_1(t)z_t + v_t \quad (16)$$

and

$$y_t = C_0(t) + C_1(t)z_t + w_t \quad (17)$$

where v_t and w_t are uncorrelated and normally distributed with variance-covariance matrices $\Sigma(t)$ and $S(t)$. Equation (16) is the *transition* equation which describes the dynamics of the basic state-variables, as in (3) in the previous section. The state-variables, z_t , are observed

⁴In this recursive verification procedure, one can also show (by straightforward application of variants of the matrix result in Lemma 1 in Appendix A) that $\Omega(t)$ is a symmetric positive definite matrix while $\Gamma(t)B_2(t)$ and $\Sigma\Gamma(t)$ are symmetric positive semi-definite matrices for $\gamma \geq 1$.

indirectly through observations of the process y_t , which is related to z_t as described by the *measurement* equation (17). As a special case, the measurement equation can be an identity equation so that the general model collapses to the set-up investigated above. However, generally the state-variables are not observed directly, and the measurement formulation even allows for “noise” (i.e. w_t) in the relationship between state-variables and observations.

For notational ease, we will in the following suppress the time dependence of the matrices $\Phi_0(t)$, $\Phi_1(t)$, $C_0(t)$, $C_1(t)$, $\Sigma(t)$ and $S(t)$ in (16) and (17). Also, to make the model realistic we will assume that the realized returns, $r_{1,t}$ and x_t , are observed without error at any time point t and, thus, assuming a special blocks-of-zeros structure to the matrices C_0 , C_1 , and S involved in the observation equation (17). These assumptions imply that the realized portfolio return and the wealth of the investor are known when making portfolio decisions at time $t = 1 \dots, T - 1$. The special structure of the matrices in the measurement equation is of no importance for the form and procedure for the solution. It can be noted that the modeling allows for unobserved basic state-variables in the state-vector, s_t (recall that $z_t = (r_{1,t}, x_t, s_t)'$), and, furthermore, that the structure of the measurement equation is not a restriction with respect to, for example, the fact that general ARMA-processes can be represented in a state-space form; see, e.g., Hamilton (chapter 13, 1994).

Let $\hat{z}_t = E_t[z_t]$ and let $\hat{z}_{t+1|t} = E_t[z_{t+1}]$. The Kalman filter then provides the following updating equations:

$$\hat{z}_{t+1|t} = \Phi_0 + \Phi_1 \hat{z}_t \quad (18)$$

$$\hat{z}_{t+1} = \hat{z}_{t+1|t} + P_{t+1|t} C_1' (C_1 P_{t+1|t} C_1' + S)^{-1} (y_{t+1} - C_0 - C_1 \hat{z}_{t+1|t}) \quad (19)$$

$$P_{t+1|t} = \Phi_1 P_t \Phi_1' + \Sigma \quad (20)$$

$$P_{t+1} = P_{t+1|t} - P_{t+1|t} C_1' (C_1 P_{t+1|t} C_1' + S)^{-1} C_1 P_{t+1|t} \quad (21)$$

The recursive system in (18)–(21) describes the updating of \hat{z} as new information arrives, and this can also be used to predict future observations of the observation vector y . This insight makes it possible to establish the likelihood function for parameter estimation of any model in state-space form (see, e.g., textbook descriptions in Harvey (1989) or Hamilton (1994)).

The relevant vector of state-variables in the following portfolio analysis is the estimate \hat{z}_t of the current position of the latent state-vector z_t (as follows from the line of arguments in this paragraph). The Kalman filter describes how to obtain \hat{z}_t as new information arrives in

the form of observations y_t and provides updating equations for obtaining the first- and second order moments of \hat{z}_{t+1} (as well as y_{t+1}) conditional on information available at time t . These moments only depend on the information available at time t through \hat{z}_t and, therefore, under normality assumptions, \hat{z}_t is a sufficient statistic with respect to the future distribution of, e.g., z_{t+i} , y_{t+i} , as well as \hat{z}_{t+i} in the economy.

For solving the model for portfolio choice, the dynamics of the relevant state-vector \hat{z} is important. We have that \hat{z}_t can be described by a VAR(1) model on the form

$$\hat{z}_{t+1} = \Phi_0 + \Phi_1 \hat{z}_t + \hat{v}_{t+1} \quad (22)$$

where the innovations \hat{v}_{t+1} are uncorrelated and normally distributed:

$$\hat{v}_{t+1} \sim \mathcal{N}(0, \hat{\Sigma}(t)),$$

and where

$$\hat{\Sigma}(t) = \text{Var}_t(\hat{v}_{t+1}) = P_{t+1|t} C_1' (C_1 P_{t+1|t} C_1' + S)^{-1} C_1 P_{t+1|t}. \quad (23)$$

The state-variable dynamics in (22) and (23) is analogous to the state-variable dynamics in (3) and (4) in the previous VAR(1)-model setting, which may be seen as the special case where state-variables are observed and $\hat{z}_t = z_t$. The VAR(1)-model in (22) and (23), however, has time-varying parameter matrices. But note that all parameter matrices are pre-determined, or deterministic, functions of time. In particular, the variance-covariance matrix $\hat{\Sigma}(t)$ is updated independently of the state-variable realizations by equation (23) and through its relation to $P_{t+1|t}$ and the Kalman filter equations (20) and (21). Hence, since the relevant variance-covariance matrices are deterministic functions of time, it is possible at time 0 to obtain the relevant last period variance-covariance matrix $\hat{\Sigma}(T-1)$, simply by recursively updating $\hat{\Sigma}(t)$ using (20) and (21), and starting with the assumed initial variance-covariance estimation error matrix P_0 on the unobservable latent state-variables. This feature is important when solving investment problems by backward dynamic programming using the Bellman equation.

The following proposition summarizes the above reasoning and procedure, and it provides the optimal portfolio choice in the general state-space model setting.

Proposition 1 *Consider the dynamic optimization problem of an investor with constant relative risk aversion and preferences given by (1) who faces investment asset dynamics given by the state-space system (16) and (17). Let \hat{z}_t be the Kalman filtered value of the basic state-variable vector at time t . The optimal portfolio policy and the value function of the problem*

are then given by

$$\alpha_t = \hat{A}_0(t) + \hat{A}_1(t)\hat{z}_t \quad (24)$$

and

$$V_t = \frac{\left(e^{\hat{B}_0(t) + (1-\gamma)\hat{B}_1(t)'\hat{z}_t + (1-\gamma)\hat{z}_t'\hat{B}_2(t)\hat{z}_t} \right) W_t^{1-\gamma} - 1}{1-\gamma} \quad (25)$$

where the matrices $\hat{A}_0(t)$, $\hat{A}_1(t)$, $\hat{B}_0(t)$, $\hat{B}_1(t)$, and $\hat{B}_2(t)$ are solutions to the difference equation system (9)-(15) with Σ substituted by $\hat{\Sigma}(t)$, as defined in (23).

Proof: See Appendix A. ■

3.1 General optimal dynamic portfolio choice

By inserting the definitions of $\hat{A}_0(t)$ and $\hat{A}_1(t)$ (i.e. the difference equation system (9)-(15) with Σ substituted by $\hat{\Sigma}(t)$) into (24), the optimal portfolio policy can alternatively be expressed as,

$$\alpha_t = \Omega(t+1)^{-1} \left[H_x \Gamma(t+1)' (\Phi_0 + \Phi_1 \hat{z}_t) + \frac{1}{2} \hat{\sigma}_x^2(t) + (1-\gamma) H_x \hat{\Sigma}(t) \Gamma(t+1) (\hat{B}_1(t+1) + H_1') \right] \quad (26)$$

with $\hat{\sigma}_x^2(t) = \text{diag}(\hat{\Sigma}_{xx}(t))$, and where $\hat{\Sigma}_{xx}(t) = H_x \hat{\Sigma}(t) H_x'$ is the excess return variance-covariance matrix given all information available at time t . In order to interpret the optimal portfolio policy in (26), it is constructive to consider two well-known special cases: (i) the case of logarithmic utility ($\gamma = 1$), and (ii) the case of a myopic investor.

In the special case of a log-investor, the last term in (26) vanishes. Moreover, by inspecting the definitions of $\Omega(t+1)$ and $\Gamma(t+1)$ in (11) and (12), it is seen that $\Omega(t+1) = \Sigma_{xx}(t)$ and that $\Gamma(t+1)$ is the identity matrix in this special case. Hence, for a log-investor the optimal portfolio policy reduces to

$$\alpha_t = \hat{\Sigma}_{xx}(t)^{-1} \left[H_x (\Phi_0 + \Phi_1 \hat{z}_t) + \frac{1}{2} \hat{\sigma}_x^2(t) \right] = \hat{\Sigma}_{xx}(t)^{-1} \left[\mathbb{E}_t[x_{t+1}] + \frac{1}{2} \hat{\sigma}_x^2(t) \right] \quad (27)$$

In establishing the last equality in (27), we have used that the matrix H_x picks out the excess return elements of the state-variable vector, z_{t+1} (and that $\hat{x}_{t+1} = x_{t+1}$ since the excess return elements of the state-variable vector are assumed perfectly observed). Since x_{t+1} is the vector of log-excess returns, the term, $\mathbb{E}_t[x_{t+1}] + \frac{1}{2} \hat{\sigma}_x^2(t)$, describes the vector of expected excess returns given the available information at time t . If all state-variables in z_t were perfectly observed, we would have that $\mathbb{E}_t[x_{t+1}] = H_x(\Phi_0 + \Phi_1 z_t)$. But in the general case, the relevant expectations

are based on the filtered (or, *perceived*) state-variable values, i.e. $\mathbb{E}_t[x_{t+1}] = H_x(\Phi_0 + \Phi_1 \hat{z}_t)$. The formula for the optimal portfolio of a log-investor (also known as the growth-optimal portfolio) is identical to similar formulas in Samuelson (1969) and Merton (1969), except that in the present case the excess returns are replaced by their perceived values and the relevant variance-covariance matrix $\hat{\Sigma}(t)$ incorporates the uncertainty with respect to the fact that some of the basic state-variables are not being perfectly observed.

A myopic investor has investment horizon $T = t + 1$ at time t , and in this special case we thus have that $\hat{B}_1(t + 1) = 0$ and $\hat{B}_2(t + 1) = 0$. Also, again by inspecting the definitions of $\Omega(t + 1)$ and $\Gamma(t + 1)$ in (11) and (12), it is seen that $\Omega(t + 1) = \gamma \Sigma_{xx}(t)$ and that $\Gamma(t + 1)$ is the identity matrix. Hence, for a myopic investor the optimal portfolio policy reduces to

$$\alpha_t = \frac{1}{\gamma} \hat{\Sigma}_{xx}(t)^{-1} [\mathbb{E}_t[x_{t+1}] + \frac{1}{2} \hat{\sigma}_x^2(t) + (1 - \gamma) \hat{\sigma}_{1x}(t)] \quad (28)$$

where $\hat{\sigma}_{1x}(t) = H_x \hat{\Sigma}(t) H_1'$ is the covariance vector between the return on benchmark asset 1 and the excess returns on the other $n - 1$ risky assets given information available at time t . Whenever asset 1 is risk-free (i.e. $\hat{\sigma}_1 = 0$ and $\hat{\sigma}_{1x} = 0$), the portfolio of risky assets coincides with the optimal portfolio of a log-investor. If the benchmark asset 1 is risky, investors with $\gamma \neq 1$ will adjust the allocation by a term $(1 - \gamma) \sigma_{1x}$. The myopic portfolio policy in (28) is identical to the myopic portfolio policy in Campbell et al. (2003) in the case where all state-variables are observed perfectly. When state-variables are not observed perfectly, the portfolio rule in (28) basically takes the same form as in one-period Bayesian portfolio models, as pioneered by Klein and Bawa (1976) and Bawa, Brown and Klein (1979).

The general portfolio policy in (26) can be interpreted as having a myopic component, which is identical to the portfolio in (28), and a hedge term, which is simply defined residually. The term involving $\hat{B}_1(t + 1)$ in (26) is a pure hedge term. The remaining terms are closely related, but not identical, to the myopic term.

4 Illustrative examples

To illustrate our framework we consider two relatively simple examples. In both cases the investor allocates wealth between a money market account yielding a constant interest rate, r , and a risky stock with an unobservable risk premium, μ_t .

In the first case μ_t is constant and the investor learns about $\mu_t = \mu$ by observing realized stock returns. We find the optimal investment strategy for an investor that recognizes that

he will learn more about μ during the time he invests in the market. This setup has been analyzed by Brennan (1998) in continuous-time.

In the second case we let μ_t be stochastic and mean-reverting, and innovations in μ_t may be correlated with innovations in realized stock returns. Again the investor learns about μ_t by observing realized stock returns. However, while he in the previous case eventually learns everything about μ_t , his learning in this setting is limited since μ_t is stochastic. We analyze the optimal investment strategy in steady state where no more learning is possible in the sense that the investor does not increase the precision on his estimate on μ_t over time. This setting has been analyzed by Rodriguez (2002) in continuous-time.

Throughout the examples and calibration analysis we use a monthly rebalancing frequency unless otherwise stated.

4.1 Asset allocation with learning about a constant mean

In the first case the dynamics of investment opportunities are given by

$$r_{1,t+1} = r \tag{29}$$

$$x_{t+1} = \mu_t + v_{x,t+1} \tag{30}$$

$$\mu_{t+1} = \mu_t, \tag{31}$$

where $v_{x,t+1} \sim \mathcal{N}(0, \sigma_x)$. Equations (29), (30), and (31) represent the transition equation system (16) in this example. The constant excess return volatility parameter, σ_x , and the constant risk-free interest rate, r , are known by the investor, whereas the constant expected excess return on the stock, $\mu_t = \mu$, is unobservable and unknown to the investor. The measurement equation is not stated explicitly here, but simply describes the situation where the risk-free returns and realized excess stock returns are observed perfectly (i.e. without measurement errors). The investor has a subjective prior on the risk premium μ_t given by a normal distribution with mean m_t and variance p_t . This distribution is subsequently updated through the Kalman filter recursions as realized excess stock returns are observed. In the continuous-time limit, the above dynamic asset allocation model with learning is equivalent to the model analyzed by Brennan (1998) who assumes a constant interest rate and that the stock price follows a geometric diffusion process with unknown drift parameter.

In the numerical implementation reported below, the prior mean and standard deviation of the monthly risk premium are given by 0.0048 and 0.0017, respectively. This corresponds

to the sample mean and standard deviation of the sample mean using monthly data on the S&P 500 stock index from March 1951 to June 2004 (which is part of the data used in the subsequent calibration analysis). The standard deviation of monthly realized log excess stock returns is set to 0.0418 and the monthly log interest rate is set to 0.0043 - both equal to their sample counterparts.

Table 1 reports the fraction of wealth allocated to stocks for varying risk-aversion coefficients, investment horizons and investor types. α denotes the allocation of a long-term investor who incorporates uncertainty about μ_t and learning into his investment decision. α^* denotes the allocation of a myopic investor who takes into account uncertainty about μ_t (but, obviously, does not take into account the effect of future learning). Finally, $\tilde{\alpha}^*$ denotes the allocation of a myopic investor who ignores uncertainty about μ_t and takes his prior mean as the true value. Uncertainty about μ_t has two effects. First, the variance of the posterior stock return distribution is increased, which decreases the stock allocation for myopic investors (compare α^* with $\tilde{\alpha}^*$). This effect dates back to Klein and Bawa (1976) and Bawa, Brown and Klein (1979). The effect disappears in continuous-time and, hence, is not present in Brennan (1998). Note also that this effect - at least for the parameter values chosen here - tends to be small (risk premium uncertainty adds only about 0.1 percent to the monthly volatility of stock returns).

The second, and more important, effect is associated with learning about the true risk premium over time. Higher-than-expected stock returns will lead to an upward assessment of the risk premium. This creates a (perfect) positive correlation between innovations to realized stock returns and innovations to the risk premium estimate. Hence, although true investment opportunities are constant, the investor perceives them as being time-varying, and in order to hedge these perceived variations the investor will decrease his allocation to stocks. As indicated in the table - and noted first by Brennan (1998) - this effect is quite significant (compare α with α^*).⁵

⁵Of course, the quantitative effects of learning in general depend critically on the assumed mean and variance of the prior distribution.

4.2 Asset allocation with unobservable mean-reverting risk premium

In the second case the dynamics of investment opportunities are given by

$$r_{1,t+1} = r \quad (32)$$

$$x_{t+1} = \mu_t + v_{x,t+1} \quad (33)$$

$$\mu_{t+1} = \bar{\mu}(1 - \phi_\mu) + \phi_\mu\mu_t + v_{\mu,t+1}, \quad (34)$$

where $v_{x,t+1} \sim \mathcal{N}(0, \sigma_x^2)$, $v_{\mu,t+1} \sim \mathcal{N}(0, \sigma_\mu^2)$ and $v_{x,t+1}$ and $v_{\mu,t+1}$ have correlation ρ . This discrete-time model of investment opportunities is similar to the infinite horizon asset allocation model considered by Campbell and Viceira (1999) in the case of full and perfect observability. Equations (32), (33), and (34) represent the transition equation system (16) in our example. Again, the measurement equation is not stated explicitly for convenience, but it is simply assumed that the risk-free returns and realized excess stock returns are observed perfectly, whereas the equity risk premium, μ_t , is unobservable. The constant risk-free interest rate, r , and the excess stock return parameters $\bar{\mu}$, ϕ_μ , σ_x , σ_μ , and ρ are all known by the investor. In particular, $\bar{\mu}$ describes the unconditional (or steady state) mean of the equity risk premium, while ϕ_μ describes the degree of mean-reversion of the equity risk premium. In this setting we focus on optimal asset allocation when no more learning is possible (in the sense that more return observations will not decrease the variance of the risk premium estimate) and, hence, the prior distribution of μ_t is of no importance. To this end, we iterate the Kalman filter recursions until $\hat{\Sigma}(t)$ converges to its steady state value $\hat{\Sigma}$ (convergence to steady state is guaranteed for any positive semi-definite initial variance-covariance matrix, provided that the system is stationary, see Hamilton (p. 390, 1994)) which then provides the basis for calculating the optimal investment strategy. To calibrate the model parameters, we follow Campbell and Viceira (1999) and assume that the risk premium is really driven by the dividend-price ratio.⁶ The calibrated parameter values applied below are given in Table 2. Table 2 also contains limiting and equivalent continuous-time model parameter estimates (see Campbell,

⁶More specifically, we first estimate the following model by maximum likelihood using monthly data from March 1951 to June 2004:

$$\begin{aligned} x_{t+1} &= c_x + \phi_x(D/P)_t + v_{x,t+1} \\ (D/P)_{t+1} &= c_{DP} + \phi_{DP}(D/P)_t + v_{DP,t+1}. \end{aligned}$$

The parameters of (34) are then obtained as: $\bar{\mu} = c_x + \phi_x c_{DP} / (1 - \phi_{DP})$, $\phi_\mu = \phi_{DP}$ and $\sigma_\mu^2 = \phi_x^2 \sigma_{DP}^2$.

Chacko, Rodriguez and Viceira (2004) and conversion formulas therein) which are used below to address the question of rebalancing frequency importance in the present context.

Table 3 reports the fraction of wealth allocated to stocks for varying risk-aversion coefficients, investment horizons and investor types when the risk-premium equals its steady state mean and with monthly rebalancing. α and α^* again denote the allocations of a long-term and a myopic investor, respectively, who do not observe the risk premium. As a benchmark we also consider the allocations of a long-term and a myopic investor - denoted $\tilde{\alpha}$ and $\tilde{\alpha}^*$, respectively - who do observe the risk premium. Again the effect of risk-premium uncertainty on the allocation of myopic investors is very small (compare α^* with $\tilde{\alpha}^*$). However, the effect on the allocation of long-term investors is pronounced. In this case the investor who observes the true risk-premium has a large positive hedge demand due to the negative correlation between innovations to the risk premium and innovations to realized returns. Since the filtered risk-premium is less variable than the true risk premium, the investor who only observes the risk premium indirectly perceives investment opportunities to be less volatile and consequently has a lower hedge demand.

Allocations are often very sensitive to variations in the risk premium. Campbell and Viceira (1999) is a case in point. In Figure 1 we show the fraction of wealth allocated to stocks over time in a simulated sample of 500 observations. We report the allocation for an investor who observes the risk premium and for one who filters it out by observing realized stock returns. The investment horizon is 10 years and $\gamma = 5$. The allocations for the investor who does not observe the risk premium are significantly less volatile than the allocations for the investor who does observe the risk premium.⁷

In Appendix B we illustrate the continuous-time equivalent to the discrete-time setup above. This allows us to get a sense of how fast our discrete-time solution converges to its continuous-time limit as the rebalancing interval goes to zero. In Table 4 we again assume an investment horizon of 10 years and $\gamma = 5$ and compute the optimal fraction of wealth allocated to stocks with discrete rebalancing - at intervals ranging from annually to daily - and continuous rebalancing. Again we report the allocation both for an investor who observes the risk premium and for one who filters it out by observing realized stock returns. The relevant parameters values for the continuous-time return dynamics in Appendix B are provided in Table 2. These are computed from the parameter estimates in the discrete-time return dynamics

⁷The average allocation to stocks is lower for the investor who does not observe the risk premium, reflecting primarily the lower average hedge demand.

in the previous section, as described above. Furthermore, the discrete-time return dynamics for various intervals are obtained by appropriately discretizing the continuous-time return dynamics (cf. Campbell, Chacko, Rodriguez and Viceira (2004)). We see that the discrete-time solution converges quite fast to the continuous-time solution as the rebalancing interval shrinks toward zero. At daily rebalancing the discrete-time and continuous-time solutions are virtually indistinguishable. At monthly rebalancing - which is what we use throughout the main part of the paper - the discrete-time solution remains very close to the continuous-time solution.

5 Asset Allocation with time-varying real interest rates and bond and stock risk premia

We will now consider and calibrate a relatively realistic model of dynamic investment opportunities, which fits into the general framework considered in Section 3. We will start out setting up the model in general terms in Section 5.1. In Section 5.2 the model is specialized and the estimation approach and results are reported. Asset allocation implications are then considered in Section 5.3

5.1 The model

Uncertainty is generated by a five-dimensional standard Wiener process, W_t , and the underlying state vector is three-dimensional and denoted X_t in the following. The dynamics of the state-vector is given by the Gaussian process:

$$dX_t = -KX_t + \sigma_X dW_t, \quad (35)$$

where the matrices K and σ_X have dimensions 3×3 and 3×5 , respectively.

The term-structure part of the model is of the essentially affine type discussed in Duffee (2002) and Dai and Singleton (2002) and shown in these papers to be capable of replicating excess returns predictability (and hence deviations from the expectations hypothesis) in the bond market. The instantaneous nominal interest rate is affine in the state-variables:

$$r(X_t) = \delta_0 + \delta_X X_t. \quad (36)$$

And the vector of market prices of risk is likewise affine in the state-variables:

$$\lambda(X_t) = \lambda_0 + \lambda_X X_t, \quad (37)$$

where λ_0 and λ_X have dimensions 3×1 and 3×3 , respectively. This implies that the instantaneous return on a zero-coupon bond, $B_t^\tau = B(X_t, \tau)$, with time to maturity τ is given by

$$\frac{dB_t^\tau}{B_t^\tau} = (r(X_t) + \sigma_B \lambda(X_t)) dt + \sigma_B dW_t, \quad (38)$$

where $\sigma_B = -A_2'(\tau)\sigma_X$ and $A_2(\tau)$ are described in Appendix C.

The instantaneous risk premium on stocks η is assumed affine in the state-variables:

$$\eta(X_t) = \eta_0 + \eta_X X_t, \quad (39)$$

which allows for excess return predictability in the stock market (we could alternatively write the risk premium in terms of market prices of risk). The return on a stock S is then:

$$\frac{dS_t}{S_t} = (r(X_t) + \eta(X_t)) dt + \sigma_S dW_t. \quad (40)$$

Finally, expected inflation i is assumed affine in the state-variables:

$$i(X_t) = \zeta_0 + \zeta_X X_t, \quad (41)$$

and the price level Π evolves according to

$$\frac{d\Pi_t}{\Pi_t} = i(X_t) dt + \sigma_\Pi dW_t. \quad (42)$$

We assume that the investor allocates wealth between a money market account (A_t), stocks and 10yr zero-coupon bonds (allowing for investment in multiple bonds with different maturities is straightforward). To map the model into the setup of section 2.2 we consider the expanded state vector $\Psi = \left(\log \frac{A_t}{\Pi_t}, \log \frac{S_t}{A_t}, \log \frac{B_t^{10}}{A_t}, X_t \right)$. The first three elements of $d\Psi$ denote the instantaneous log real return, the instantaneous log excess stock return and the instantaneous log excess return on the 10yr bond. The dynamics of Ψ is given by the Gaussian vector process

$$d\Psi = (\varphi_\Psi - K_\Psi X_t) dt + \sigma_\Psi dW_t, \quad (43)$$

where

$$\varphi_\Psi = \begin{pmatrix} \delta_0 - \zeta_0 + \frac{1}{2}\sigma_\Pi^2 \\ \eta_0 - \frac{1}{2}\sigma_S^2 \\ \sigma_{B(10)}\lambda_0 - \frac{1}{2}\sigma_{B(10)}^2 \\ 0_{3 \times 1} \end{pmatrix}, K_\Psi = \begin{pmatrix} 0_{1 \times 3} & \delta_X - \zeta_X \\ 0_{1 \times 3} & \eta_X \\ 0_{1 \times 3} & \sigma_{B(10)}\lambda_X \\ 0_{3 \times 3} & -K \end{pmatrix}, \sigma_\Psi = \begin{pmatrix} -\sigma_\Pi \\ \sigma_S \\ \sigma_{B(10)} \\ \sigma_X \end{pmatrix}. \quad (44)$$

Discretizing a Gaussian process is straightforward (see, e.g., Karatzas and Shreve (pp. 354-357, 1991), and in this way we obtain the VAR(1) for $z_t = (r_{1,t}, x_t, s_t)$ given in section 2.2, where x_t contains the log excess stock and bond return and s_t equals X_t .

The investor observes perfectly the realized 1mth real interest rate and the realized 1mth excess returns on 10yr zero-coupon bonds and stocks (the first three elements of the z_t -vector). In addition, the investor observes with measurement errors a term structure of nominal zero-coupon bond yields (with maturities of 1mth, 6mth, 1yr, 2yr, 5yr and 10yr) and the dividend-price ratio. Appendix C shows that bond yields are affine in the state-variables:

$$y(X_t, \tau) = \frac{A_1(\tau)}{\tau} + \frac{A_2(\tau)}{\tau} X_t. \quad (45)$$

Equation (45), with “measurement noise” added, thus describes one of the entries in the measurement equation. This way of using panel data in combination with the Kalman filter for estimating term-structure models was pioneered by Pennacchi (1991) and has been applied in numerous papers since.

Appendix C also demonstrates that, using the log-linear dividend discount model of Campbell and Shiller (1988a,1988b), the log dividend-price ratio is approximately affine in the state variables:

$$d_t - p_t = -\frac{k + g + \frac{1}{2}\sigma_S^2 - \delta_0 - \eta_0}{1 - \rho} + (\delta_X + \eta_X)\nu(I - \rho\Phi)^{-1}X_t. \quad (46)$$

Equation (46), with “measurement noise” added, likewise describes an entry in our measurement equation. Brennan and Xia (2004) have recently argued that using a dividend discount model is the preferable way to estimate the - otherwise hard to detect - process for expected stock returns (they do not decompose expected stock returns into interest rates and stock risk premia). They employ what is essentially a non-linear Kalman filter to estimate the expected stock return process. The approach in this paper represents a simpler - but approximate - alternative to theirs. On the other hand, our model is richer than theirs in the sense that we model interest rates and stock risk premia separately and both interest rates and stock risk premia are allowed to depend on multiple state-variables. To simplify the analysis we have assumed that expected future dividend-growth is constant. An extension to time-variations in expected future dividend growth is straightforward, but Brennan and Xia (2004) find that such an extension does not change the results regarding expected stock returns much.

The measurement errors in the bond yield and dividend-price ratio relations are assumed to be mean-zero normally distributed and serially and cross-sectionally uncorrelated. In addition we assume that the percentage standard deviation of the pricing errors are equal across the

bond yield and log dividend-price ratio relations - a parametrization that seems plausible and reduces the parameter space.

5.2 Estimation results

We use monthly US data from the Treasury-Federal Reserve Accord in March 1951 to June 2004. Zero-coupon bond yields with maturities 1mth, 6mth, 1yr, 2yr, 5yr and 10yr for the period March 1951 to February 1991 are from McCulloch (1990) and McCulloch and Kwon (1993). The yields for the period March 1991 to June 2004 are constructed with bootstrapping methods from the constant maturity yields reported in the Federal Reserve H.15 statistical release. All yields are recorded end-of-month. The cum dividend stock index is constructed from end-of-month values of the S&P 500 stock index. The dividends that enter in the dividend-price ratio are calculated as the sum of dividends payed out over the past year. The consumer price index is the CPI-All Urban Consumers, All Items index published by the U.S. Bureau of Labor Statistics.

Due to the proliferation of parameters, we need to impose some structure on the model. Guided by the existing term structure literature, which shows that a two factor model can capture much of the term structure dynamics, we let the short-term interest rate be a function of the first two state-variables. Also, since the slope of the yield curve has been shown to be the most potent predictor of excess bond returns and since the second factor in term structure models usually have a high correlation with the slope, we let the market price of risk be a function of the first two state-variables. In order for the model to be econometrically identified, K must be lower triangular and $\sigma_X = (I_{3 \times 3} 0_{3 \times 2})$ (see Dai and Singleton (2000)). It must also be the case that σ_S only loads on the first four Wiener processes while σ_Π can load on all five.

Maximum likelihood estimates and asymptotic heteroscedasticity-consistent standard errors are given in pro-anno terms in Table 5 and Table 6. The mean-reversion parameters indicate that the state-variables are quite persistent. Shocks to each of the state-variables have half-lives of 5.0 years, 1.4 years and 20.1 years, respectively. To ease interpretation of the table, note that the first state-variable is highly correlated with the interest rate level (correlation with the 3mth rate is 0.981), the second state-variable is negatively correlated with the slope of the yield curve (correlation with the 10yr - 3mth spread is -0.488) and the third state-variable is highly negatively correlated with the log dividend-price ratio (correlation is -0.961).

The estimate of the η_X -vector reveals that the stock risk premium depend negatively on all three state-variables. Given the correlations above, this is consistent with the dividend-yield predicting excess stock returns with a positive sign and the short rate predicting excess stock returns with a negative sign (and more weakly, the slope of the yield curve predicting excess stock returns with a positive sign) - consistent with what has been documented in the vast return predictability literature. The estimate of the λ_X -matrix is harder to interpret. However, a simulation exercise along the line of Dai and Singleton (2002) shows that the estimated model is able to replicate the predictability of excess bond returns (across the entire maturity spectrum) from the slope of the yield curve - a widely documented fact in the literature on tests of the expectations hypothesis.⁸

Table 7 compares unconditional sample and implied moments of central variables of interest. It first reports the sample mean and standard deviation of the real interest rate and the sample means, standard deviations and Sharpe ratios of the 1yr, 5yr and 10yr excess bond returns and the excess stock return.⁹ It then reports the corresponding numbers implied by the model. The model fits the volatilities of all variables very well. The fit of the means is less impressive, although - as we explain - this is entirely expected. The implied 10yr bond risk premium is about 75bp higher than in the data. This, however, is partly due to the fact that the 10yr yield is higher toward the end than in the beginning of the sample period. This has depressed realized returns in the sample and has caused the sample mean to be lower than the true population mean. A similar explanation underlies the large discrepancy between the implied equity risk premium of approximately 4.8 percent and the sample equity risk premium of approximately 6.8 percent. In the sample there is a downward trend in the dividend-price ratio. Our model attributes all of this decline to a decline in discount rates (since we assume constant expected dividend-growth) which has inflated realized returns in the sample and caused sample mean returns to exceed the true population mean. This underscores an important advantage of our model: By utilizing the present-value relation for equities, our risk-premium estimate is not distorted by the extremely high returns in the sample. As a consequence of the above the difference between implied Sharpe ratios on stocks and bonds is much lower than the difference between sample Sharpe ratios. As we demonstrate below, this helps solve one problem that plagues many asset allocation models: That the optimal unconditional allocation almost always involves a highly leveraged position in stocks.

⁸The simulation results are available from the authors upon request

⁹Means of log returns are added one-half their variances to account for Jensen's inequality.

Figure 2 shows the fitted values for the instantaneous expected real interest rate and the instantaneous expected total and excess returns on 10yr bonds and stocks. The fitted values are calculated on the basis of the Kalman filtered smoothed state-variables. The fitted expected excess bond return corresponds to that shown in Duffee (2002) for his $A_0(3)$ model.¹⁰ The risk premium is low and stable in the 1950s and 60s, more volatile and generally higher in the 1970s and 80s, and has trended downwards since the mid-1980s.¹¹ The fitted expected total stock return follows closely the similar graphed fits in Brennan and Xia (2004). It declines from a high level during the 1950s, is low in the 1960s, rises in the 1970s and declines from the early 1980s until the peak of the stock-market in 2000, where it reaches negative values. Since then it has recovered slightly but remains very low by historical standards. As explained above, our model allows us to focus explicitly on the stock risk premium. The figure shows that the decline in expected stock returns during the 1980s was entirely due to declining interest rates. In contrast, the decline in expected stock returns during the 1990s was entirely due to a declining risk premium. Since 2000 the risk premium has increased rather dramatically.

Based on the above, we conclude that the model - despite being quite parsimonious - gives a realistic description of the dynamics of investment opportunities faced by investors; in the next section we solve for the optimal asset allocation on the basis of this model.¹²

5.3 Asset Allocation

5.3.1 General properties of the solution

Solving for the optimal asset allocation is straightforward using our solution algorithm of Section 3. The optimal asset allocation depends - apart from the parameter estimates - on the investment horizon, the risk aversion, and the position of the state vector. To get a sense of the asset allocation recommendations of the model, Panel A in Figure 3 shows the optimal asset allocation as a function of the investment horizon, assuming $\gamma = 5$ and that the state vector equals its unconditional mean. For this particular investor the allocation to stocks and bonds

¹⁰See Duffee (2002), Figure 1, Panel F.

¹¹Duffee's estimated risk premium is more volatile due to his use of a three-factor model, where the third factor - which is omitted in our model - is highly transitory

¹²One element that the model does not capture is variations in second moments. The evidence in Chacko and Viceira (2004) seems to indicate that variations in volatilities are not very important for long-term asset allocations since shocks to volatilities are relatively short-lived. Shocks to correlations seem to be more persistent and could be more important.

increases with the horizon. The stock-bond ratio also increases with the horizon, especially for investment horizons above 10 years.

The horizon effect is due to the investor's desire to hedge variations in the real short-term interest rate and bond and stock risk premia. To ease interpretation of the hedge portfolio we consider three cases in which the investor hedges only one dimension of investment opportunities (either the real interest rate, the bond risk premium or the stock risk premium) while treating the other parts as constants and equal to their unconditional means. These results are reported in panel B-D in Figure 3. The results can be understood by relating to Table 8 which shows the correlation between innovations to the three state-variables and innovations to realized bond and stock returns, the expected real interest rate and bond and stock risk premia. Innovations to the realized return on the 10yr bond is correlated with innovations to the first two state-variables (correlations are strongly negative and moderately positive, respectively), and innovations to realized stock return is strongly positively correlated with innovations to the third state-variable. Hence, the 10yr bond can be used to hedge exposure to the first two state-variables while the stock can be used to hedge exposure to the third state-variable

Figure 3.B shows that the hedge portfolio associated with real interest rate risk consists almost entirely of a long position in bonds. The reason is that the real interest rate depends mostly and with a positive sign on the first state-variable. Since innovations to this state-variable is strongly negatively correlated with realized bond returns, it follows that bonds provide a good hedge against time-varying real interest rates. That nominal bonds provide a better hedge against real interest rate risk than stocks is not uncontroversial and somewhat model and data dependent. Brennan and Xia (2002) and Sangvinatsos and Wachter (2004) finds a similar result, but Campbell and Viceira (2002) conclude that stocks provide the better hedge.¹³

Next we consider the hedge portfolio associated with expected bond return risk. Figure 3.C shows that it also almost exclusively consists of a long bond position. The explanation is that bond risk premia depend mostly and negatively on the second state-variable. Since innovations to this state-variable is moderately positively correlated with realized bond returns, it follows that bonds also provide a good hedge for time-varying bond risk premia. Sangvinatsos and Wachter (2004) reach a similar conclusion.

¹³That conclusion is sensitive to the period over which their model is estimated. When they estimate their model on data from 1982-1999, they find that nominal bonds provide the best hedge.

The hedge portfolio associated with expected stock return risk is given in Figure 3.D, which shows that it consists of a long stock position and a short bond position. Stock risk premia depends negatively on all three state-variables so - again from the correlation structure in Table 8 - it follows that a portfolio short the 10yr bond and long the stock will provide a good hedge against time-varying stock risk premia. The role of stocks in hedging time-varying stock risk premia has been stressed in quite a few papers and also emerged from the analysis in Section 4.2 (see the references there).

In Figure 4 we analyze the optimal asset allocation as a function of risk aversion in the same four cases discussed above. We assume an investment horizon of 10 years and that the state vector equals its unconditional mean. The structure of the figure is similar to the structure of Figure 3. We see that above $\gamma = 2$, the composition of the hedge portfolio is not very sensitive to reasonable variations in risk aversion. The myopic portfolio, of course, is very sensitive to the risk aversion level.

5.3.2 The impact of measurement errors

In the estimation and asset allocation above we have assumed that the measurement errors in the zero-coupon bond yield relations and the log dividend-price ratio relation have equal standard deviations. We now investigate the impact on the stock allocation of varying the measurement error in the dividend yield relation. This can be thought of as varying the degree of confidence the investor has in this relation. We consider two cases. In the first case the investor has full confidence and assumes zero measurement error when estimating the parameters and subsequently solving the asset allocation problem. In the second case the investor is quite skeptic and assumes that the percentage standard deviation of the dividend yield pricing error is twice the percentage standard deviation of the bond yield pricing errors. Figure 5 shows the optimal allocation to stocks as a fraction of wealth over time in these two cases and the base-line case from above. We assume an investment horizon of ten years and relative risk aversion of five. The allocations in the three cases follow the same overall trend but the less confidence the investor has in the dividend-price ratio relation, the less volatile are the allocations. The standard deviations of the allocations are 86.1 percent in the high confidence case, 52.3 percent in the moderate confidence case and 43.7 percent in the low confidence case.¹⁴

¹⁴Surprisingly the allocations in the case of low confidence is relatively volatile during the monetary experiment in the early 1980's. The reason is that stock risk premium in this case is estimated to depend relatively

6 Conclusion

This paper has provided a basic framework for addressing the importance of time horizon and risk aversion on optimal portfolio choice in VAR-models. The modeling allows for state-variables being unobservable, and examples have demonstrated that this significantly affects dynamic portfolio allocations. In particular, the investor reacts more conservatively to news since portfolio allocations are less variable over time compared to the case where state-variables are assumed directly and perfectly observable.

We conjecture that the current research can be extended in various direction. It will, for example, be possible to incorporate consumption and possibly other features by application of the approximate solution technique suggested and outlined by Campbell and Viceira (2002) and Campbell et al. (2003). Furthermore, the general state-space formulation is rich and allows for representation of many alternative return generating models which have not been addressed in this paper.

more strongly on the term structure factors which are more volatile during this period.

Appendix A

The proof of Proposition 1 as well as the solution forms in (7), (8), (24), and (25) and the recursive solution algorithm in (9)–(15) hinges on Lemma 3 below. Lemma 3 on the other hand is based on the following Lemma 2, while the matrix result provided in Lemma 1 is relevant for showing that the matrices obtained by the recursive solution algorithm are well-defined. The result in Lemma 2 can be obtained by basically rewriting and elaborating on a similar result in Campbell, Chan and Viceira (pp. 10-11, 2002), (under slightly different notation). An explicit proof of Lemma 2, however, is stated for completeness and for direct use and adoption in the subsequent proofs.

Lemma 1 *Let A and B be symmetric positive semi-definite matrices. Then*

$$C = (I + BA)^{-1}B$$

is well-defined and symmetric positive semi-definite.

Proof: We will first prove that $I + BA$ has full rank. To this end, we will start by assuming that A is positive definite and thus invertible. Now assume that $I + BA$ does not have full rank. It is then possible to find a vector $v \neq 0$ so that

$$v'(I + BA) = 0 \Rightarrow v' = -v'BA \Rightarrow v'A^{-1} = -v'B \Rightarrow v'A^{-1}v = -v'Bv.$$

But since A^{-1} is positive definite and B is positive semi-definite, the last equality can only be valid if $v = 0$; this contradicts that $v \neq 0$ and, hence, $I + BA$ must have full rank whenever A is strictly positive definite. In the general case, let $\{A_n\}$ be a sequence of symmetric positive definite matrices such that $A_n \rightarrow A$ as $n \rightarrow \infty$. Again, assume that $I + BA = \lim_{n \rightarrow \infty} (I - BA_n)$ does not have full rank. Then it is possible to find a vector $v \neq 0$ so that $\lim_{n \rightarrow \infty} v'(I - BA_n) = 0$. Reestablishing the above line of implications, this can only be the case if

$$\lim_{n \rightarrow \infty} v'A_n^{-1}v = -v'Bv.$$

Since the right-hand side is non-positive and $v'A_n^{-1}v$ is strictly positive for all n , this can only be valid if $v = 0$; the proof that $I + BA$ has full rank therefore again follows by contradiction. It follows that $(I + BA)^{-1}$ and C are well-defined.

In order to prove that C is symmetric positive semi-definite, let $\{B_n\}$ be a sequence of symmetric positive definite matrices such that $B_n \rightarrow B$ as $n \rightarrow \infty$. Then

$$C_n = (I + B_nA)^{-1}B_n = (B_n^{-1}(I + B_nA))^{-1} = (B_n^{-1} + A)^{-1}$$

defines a sequence of symmetric positive definite matrices. Since C is well-defined, and $C_n \rightarrow C$ as $n \rightarrow \infty$, C will be symmetric positive semi-definite. \blacksquare

Lemma 2 *Let \tilde{z} have a multivariate normal distribution, $\tilde{z} \sim \mathcal{N}(\mu_z, \Sigma_{zz})$, and let B_0 be a constant, B_1 an n -dimensional vector, and B_2 a symmetric and negative semi-definite $n \times n$ matrix. Then*

$$\mathbb{E}_t \left[e^{B_0 + B_1' \tilde{z} + \tilde{z}' B_2 \tilde{z}} \right] = e^{C_0 + C_1' \mu_z + \mu_z' C_2 \mu_z}$$

where

$$C_0 = B_0 + \frac{1}{2} \ln |\Gamma| + \frac{1}{2} B_1' \Sigma_{zz} \Gamma B_1$$

$$C_1 = \Gamma B_1$$

$$C_2 = \Gamma B_2$$

$$\Gamma = (I - 2B_2 \Sigma_{zz})^{-1}$$

Furthermore, C_2 is a symmetric and negative semi-definite matrix.

Proof: Let $\Psi = (\Sigma_{zz}^{-1} - 2B_2)^{-1}$. Since Σ_{zz}^{-1} is symmetric and positive definite and $-2B_2$ is symmetric and positive semi-definite, we have that Ψ is symmetric and positive definite (and, thus, invertible). Using the definition, invertibility, and symmetry of Ψ , we obtain

$$\begin{aligned} \mathbb{E}_t \left[e^{B_0 + B_1' \tilde{z} + \tilde{z}' B_2 \tilde{z}} \right] &= \int_{\mathbb{R}^n} e^{B_0 + B_1' z + z' B_2 z} |\Sigma_{zz}|^{-\frac{1}{2}} (2\pi)^{-\frac{n}{2}} e^{-\frac{1}{2} (z - \mu_z)' \Sigma_{zz}^{-1} (z - \mu_z)} dz \\ &= \int_{\mathbb{R}^n} |\Sigma_{zz}|^{-\frac{1}{2}} (2\pi)^{-\frac{n}{2}} e^{-\frac{1}{2} z' (\Sigma_{zz}^{-1} - 2B_2) z + (B_1' + \mu_z' \Sigma_{zz}^{-1}) z + B_0 - \frac{1}{2} \mu_z' \Sigma_{zz}^{-1} \mu_z} dz \\ &= \int_{\mathbb{R}^n} |\Sigma_{zz}|^{-\frac{1}{2}} (2\pi)^{-\frac{n}{2}} e^{-\frac{1}{2} z' \Psi^{-1} z + (B_1' + \mu_z' \Sigma_{zz}^{-1}) z + B_0 - \frac{1}{2} \mu_z' \Sigma_{zz}^{-1} \mu_z} dz \\ &= \int_{\mathbb{R}^n} |\Psi|^{-\frac{1}{2}} (2\pi)^{-\frac{n}{2}} e^{-\frac{1}{2} (z - \Psi(B_1 + \Sigma_{zz}^{-1} \mu_z))' \Psi^{-1} (z - \Psi(B_1 + \Sigma_{zz}^{-1} \mu_z))} dz \times \\ &\quad |\Psi|^{\frac{1}{2}} |\Sigma_{zz}|^{-\frac{1}{2}} e^{\frac{1}{2} (B_1 + \Sigma_{zz}^{-1} \mu_z)' \Psi (B_1 + \Sigma_{zz}^{-1} \mu_z) + B_0 - \frac{1}{2} \mu_z' \Sigma_{zz}^{-1} \mu_z}, \end{aligned}$$

where the term inside the last integral is the density function for an $N(\Psi(B_1 + \Sigma_{zz}^{-1} \mu_z), \Psi)$ distributed random variable which by definition integrates to one. Hence we have

$$\begin{aligned} \mathbb{E}_t \left[e^{B_0 + B_1' \tilde{z} + \tilde{z}' B_2 \tilde{z}} \right] &= e^{\frac{1}{2} \ln |\Psi| - \frac{1}{2} |\Sigma_{zz}| + \frac{1}{2} (B_1 + \Sigma_{zz}^{-1} \mu_z)' \Psi (B_1 + \Sigma_{zz}^{-1} \mu_z) + B_0 - \frac{1}{2} \mu_z' \Sigma_{zz}^{-1} \mu_z} \\ &= e^{B_0 + \frac{1}{2} \ln |\Psi| - \frac{1}{2} |\Sigma_{zz}| + \frac{1}{2} B_1' \Psi B_1 + B_1' \Psi \Sigma_{zz}^{-1} \mu_z + \frac{1}{2} \mu_z' (\Sigma_{zz}^{-1} \Psi \Sigma_{zz}^{-1} - \Sigma_{zz}^{-1}) \mu_z} \\ &= e^{C_0 + C_1' \mu_z + \mu_z' C_2 \mu_z}, \end{aligned}$$

where the third equality follows from observing that $\Psi = \Sigma_{zz} (I - 2B_2 \Sigma_{zz})^{-1} = \Sigma_{zz} \Gamma$ and that $\Psi = (\Sigma_{zz}^{-1} - 2B_2)^{-1} \Rightarrow \Sigma_{zz}^{-1} = \Psi^{-1} + 2B_2 \Rightarrow \Sigma_{zz}^{-1} \Psi \Sigma_{zz}^{-1} - \Sigma_{zz}^{-1} = \Sigma_{zz}^{-1} (I + 2\Psi B_2 - I) = 2\Psi B_2$.

Finally, since $C_2 = \Gamma B_2 = (I - 2B_2 \Sigma_{zz})^{-1} B_2$, it follows by a straightforward application of Lemma 1 that C_2 is a symmetric and negative semi-definite matrix. \blacksquare

Lemma 3 *Let z_{t+1} be described by a VAR(1) model,*

$$z_{t+1} = \Phi_0 + \Phi_1 z_t + v_{t+1}, \quad v_{t+1} \sim \mathcal{N}(0, \Sigma).$$

If B_2 is symmetric and negative semi-definite, then

$$\mathbb{E}_t \left[e^{B_0 + B_1' z_{t+1} + z_{t+1}' B_2 z_{t+1}} \right] = e^{D_0 + D_1' z_t + z_t' D_2 z_t}$$

where

$$\begin{aligned} D_0 &= B_0 + \frac{1}{2} \ln |\Gamma| + \frac{1}{2} B_1' \Sigma \Gamma B_1 + \Phi_0' \Gamma B_1 + \Phi_0' \Gamma B_2 \Phi_0 \\ D_1 &= \Phi_1' \Gamma B_1 + 2\Phi_1' \Gamma B_2 \Phi_0 \\ D_2 &= \Phi_1' \Gamma B_2 \Phi_1 \\ \Gamma &= (I - 2B_2 \Sigma)^{-1}, \end{aligned}$$

and where D_2 is symmetric and negative semi-definite.

Proof: We have that $z_{t+1} \sim \mathcal{N}(\mu_z, \Sigma_{zz})$ where $\mu_z = \Phi_0 + \Phi_1 z_t$ and $\Sigma_{zz} = \Sigma$. By application of Lemma 2, one thus obtains

$$\begin{aligned} \mathbb{E}_t \left[e^{B_0 + B_1' z_{t+1} + z_{t+1}' B_2 z_{t+1}} \right] &= e^{C_0 + C_1' \mu_z + \mu_z' C_2 \mu_z} \\ &= e^{C_0 + C_1' (\Phi_0 + \Phi_1 z_t) + (\Phi_0 + \Phi_1 z_t)' C_2 (\Phi_0 + \Phi_1 z_t)} \\ &= e^{D_0 + D_1' z_t + z_t' D_2 z_t} \end{aligned}$$

where C_0 , C_1 , and C_2 are as stated in Lemma 2, and where the last equality is obtained by collection of constant, linear, and quadratic terms of z_t in the exponential. Thus,

$$\begin{aligned} D_0 &= C_0 + C_1' \Phi_0 + \Phi_0' C_2 \Phi_0 \\ D_1 &= C_1' \Phi_1 + 2\Phi_0' C_2 \Phi_1 \\ D_2 &= \Phi_1' C_2 \Phi_1. \end{aligned}$$

By inserting C_0 , C_1 , and C_2 , as stated in Lemma 2, one obtains D_0 , D_1 , and D_2 in the form stated in the lemma and with $\Gamma = (I - 2B_2 \Sigma)^{-1}$.

Finally, since $C_2 = \Gamma B_2$ is symmetric and negative semi-definite (cf. Lemma 2), it follows that $D_2 = \Phi_1' C_2 \Phi_1$ is also symmetric and negative semi-definite. \blacksquare

We can now prove Proposition 1 as well as the solution forms in (7), (8), (24), and (25) and the recursive solution algorithm in (9)–(15). (The notation from section 2 is used.)

Proof: (of Proposition 1) Using the conjectured form of the indirect utility function, and that next period wealth is given by $W_{t+1} = W_t e^{r_{p,t+1}}$ (and with log portfolio return $r_{p,t+1}$ as given in (5)), we can evaluate the object function in the maximization involved in the relevant Bellman equation (cf. (6)),

$$\begin{aligned}
G(t, \alpha_t) &= \mathbb{E}_t[V_{t+1}] \\
&= \mathbb{E}_t \left[\frac{\left(e^{B_0(t+1)+(1-\gamma)B_1(t+1)'z_{t+1}+(1-\gamma)z'_{t+1}B_2(t+1)z_{t+1}} \right) W_{t+1}^{1-\gamma} - 1}{1-\gamma} \right] \\
&= \mathbb{E}_t \left[\frac{\left(e^{B_0(t+1)+(1-\gamma)B_1(t+1)'z_{t+1}+(1-\gamma)z'_{t+1}B_2(t+1)z_{t+1}+(1-\gamma)r_{p,t+1}} \right) W_t^{1-\gamma} - 1}{1-\gamma} \right] \\
&= \frac{g(t, \alpha_t) W_t^{1-\gamma} - 1}{1-\gamma}
\end{aligned} \tag{47}$$

where

$$\begin{aligned}
g(t, \alpha_t) &= \mathbb{E}_t \left[e^{B_0(t+1)+(1-\gamma)B_1(t+1)'z_{t+1}+(1-\gamma)z'_{t+1}B_2(t+1)z_{t+1}+(1-\gamma)r_{p,t+1}} \right] \\
&= \mathbb{E}_t \left[e^{B_0(t+1)+(1-\gamma)B_1(t+1)'z_{t+1}+(1-\gamma)z'_{t+1}B_2(t+1)z_{t+1}+(1-\gamma)(r_{1,t+1}+\alpha'_t x_{t+1}+\frac{1}{2}\alpha'_t(\sigma_x^2-\Sigma_{xx}\alpha_t))} \right] \\
&= \mathbb{E}_t \left[\exp \left\{ (B_0(t+1) + (1-\gamma)\frac{1}{2}\alpha'_t(\sigma_x^2 - \Sigma_{xx}\alpha_t)) \right. \right. \\
&\quad \left. \left. + (1-\gamma)(B_1(t+1)' + H_1 + \alpha'_t H_x)z_{t+1} + z'_{t+1}(1-\gamma)B_2(t+1)z_{t+1} \right\} \right]
\end{aligned} \tag{48}$$

and where H_1 and H_x are selection matrices that select respectively the first element and the excess returns from z_{t+1} , i.e. $r_{1,t+1} = H_1 z_{t+1}$ and $x_{t+1} = H_x z_{t+1}$.

Using Lemma 3, $g(t, \alpha_t)$ can now be evaluated and written as

$$g(t, \alpha_t) = e^{D_0(t, \alpha_t) + D_1(t, \alpha_t)'z_t + z'_t D_2(t)z_t} \tag{49}$$

where

$$\begin{aligned}
D_0(t, \alpha_t) &= B_0(t+1) + \frac{1}{2} \ln |\Gamma(t+1)| + \frac{1}{2}(1-\gamma)\alpha_t'(\sigma_x^2 - \Sigma_{xx}\alpha_t) \\
&\quad + \frac{1}{2}(1-\gamma)^2 (B_1(t+1)' + H_1 + \alpha_t' H_x) \Sigma \Gamma(t+1) (B_1(t+1) + H_1' + H_x' \alpha_t) \\
&\quad + (1-\gamma)\Phi_0' \Gamma(t+1) (B_1(t+1) + H_1' + H_x' \alpha_t) + (1-\gamma)\Phi_0' \Gamma(t+1) B_2(t+1) \Phi_0 \\
D_1(t, \alpha_t) &= (1-\gamma)\Phi_1' \Gamma(t+1) (B_1(t+1) + H_1' + H_x' \alpha_t) + 2(1-\gamma)\Phi_1' \Gamma(t+1) B_2(t+1) \Phi_0 \\
D_2(t) &= (1-\gamma)\Phi_1' \Gamma(t+1) B_2(t+1) \Phi_1 \\
\Gamma(t+1) &= (I + 2(\gamma-1)B_2(t+1)\Sigma)^{-1}
\end{aligned}$$

The optimal portfolio weight α_t must be determined by maximizing expected next period indirect utility, i.e. by maximizing the expression $G(t, \alpha_t)$ (and with $g(t, \alpha_t)$ given in (49)) with respect to α_t . The first order conditions from the optimization problem are given by

$$\begin{aligned}
&\frac{\partial G}{\partial \alpha_t} = 0 \\
&\Downarrow \\
&\frac{\partial D_0}{\partial \alpha_t} + \frac{\partial D_1}{\partial \alpha_t} z_t = 0 \\
&\Downarrow \\
&\frac{1}{2}\sigma_x^2 + (1-\gamma)H_x \Sigma \Gamma(t+1) (B_1(t+1) + H_1') + H_x \Gamma(t+1)' \Phi_0 + H_x \Gamma(t+1)' \Phi_1 z_t \\
&\quad - (\Sigma_{xx} - (1-\gamma)H_x \Sigma \Gamma(t+1)H_x') \alpha_t = 0 \\
&\Downarrow \\
&\alpha_t = A_0(t) + A_1(t)z_t
\end{aligned}$$

where $A_0(t)$ and $A_1(t)$ are given in (9) and (10). The indirect utility function at time t can now be determined by substituting the optimal portfolio weights into the expression in (47), i.e. $V_t = G(t, A_0(t) + A_1(t)z_t)$. By evaluating this expression for V_t , one obtains

$$\begin{aligned}
V_t &= \frac{\left(e^{D_0(t, A_0(t) + A_1(t)z_t) + D_1(t, A_0(t) + A_1(t)z_t)' z_t + z_t' D_2(t) z_t} \right) W_t^{1-\gamma} - 1}{1-\gamma} \\
&= \frac{\left(e^{B_0(t) + (1-\gamma)B_1(t)' z_t + (1-\gamma)z_t' B_2(t) z_t} \right) W_t^{1-\gamma} - 1}{1-\gamma}
\end{aligned}$$

where the last equality is obtained by collection of constant, linear, and quadratic terms of z_t , and $B_0(t)$, $B_1(t)$, and $B_2(t)$ are given in (13), (14), and (15). The fact that $B_2(t)$ is symmetric and positive semi-definite follows from the discussion after equations (13), (14), and (15) in section 2. ■

Appendix B

In this appendix we present and derive the continuous-time equivalent of the discrete-time setup in section 4.2. The investment opportunity dynamics are given by

$$\frac{dA_t}{A_t} = r dt \quad (50)$$

$$\frac{dS_t}{S_t} = (r + \eta_t)dt + \sigma_S dZ_{S,t} \quad (51)$$

$$d\eta_t = \kappa(\theta - \eta_t)dt + \sigma_\eta dZ_{\eta,t}, \quad (52)$$

where $dZ_{S,t}$ and $dZ_{\eta,t}$ denote Wiener processes with correlation ρ . The risk premium is unobservable and must be inferred from the realized stock returns. Applying the Kalman-Bucy filter (see e.g. Lipster and Shiryaev (2001), chapter 12), we can derive the investment opportunity dynamics perceived by the investor:

$$\frac{dA_t}{A_t} = r dt \quad (53)$$

$$\frac{dS_t}{S_t} = (r + \hat{\eta}_t)dt + \sigma_S d\hat{Z}_{S,t} \quad (54)$$

$$d\hat{\eta}_t = \kappa(\theta - \hat{\eta}_t)dt + \sigma_{\hat{\eta}} d\hat{Z}_{S,t}, \quad (55)$$

where $\hat{\eta}_t$ denotes the estimate of the unobserved risk premium. $\sigma_{\hat{\eta}}$ and $d\hat{Z}_{S,t}$ are given by

$$\sigma_{\hat{\eta}} = \frac{\rho\sigma_\eta\sigma_S + v_t}{\sigma_S} \quad (56)$$

$$d\hat{Z}_{S,t} = \left(\frac{\eta_t - \hat{\eta}_t}{\sigma_S} \right) dt + dZ_{S,t} \quad (57)$$

and v_t - the variance of the risk premium estimate - follows

$$\frac{dv_t}{v_t} = -2\kappa v_t + \sigma_\eta^2 - \left(\frac{\rho\sigma_\eta\sigma_S + v_t}{\sigma_S} \right)^2. \quad (58)$$

Note that although the market is incomplete, an investor with the assumed information set will perceive the market as being complete. This ‘‘observational completeness’’ result has been stressed by Rodriguez (2002).

As in section 4.2 we focus on the steady state where no more learning is possible. In steady state we must have $\frac{dv_t}{dt} = 0$ which yields a quadratic equation, the non-negative root of which can be shown to equal

$$v^* = -\kappa\sigma_S^2 - \rho\sigma_\eta\sigma_S + \sigma_S \sqrt{\kappa^2\sigma_S^2 + \sigma_\eta^2 + 2\kappa\rho\sigma_\eta\sigma_S}. \quad (59)$$

Hence, in steady state the diffusion term $\sigma_{\hat{\eta}}$ is given by (56) with v_t replaced by v^* .¹⁵

The closed form solution to the optimal asset allocation strategy of a perfectly informed long-horizon CRRA investor facing investment opportunity dynamics given by (50)-(52) was derived by Kim and Omberg (1996). The same solution can be used to solve the optimal asset allocation strategy of an imperfectly informed investor facing perceived investment opportunity dynamics given by (53)-(55).

¹⁵It follows from (59) that $v^* = 0$ when $\rho = 1$ or $\rho = -1$ (the latter only holds true provided that $\kappa\sigma_S^2 - \sigma_\eta\sigma_S \geq 0$). From (56) it then follows that the true and filtered risk premium coincide in these special cases.

Appendix C

In this appendix we show that the zero-coupon bond yield is affine in the state-variables, X_t , and the log dividend-price ratio similarly is approximately affine as a function of X_t .

Given the specification of the market price of risk, the dynamics of the state-vector under the equivalent martingale measure Q is given by a Gaussian process:

$$dX_t = \tilde{K}(\tilde{\theta} - X_t) + \sigma_X d\tilde{W}_t, \quad (60)$$

where $\tilde{K} = K + \sigma_X \lambda_X$, $\tilde{\theta} = -\tilde{K}^{-1} \sigma_X \lambda_0$ and \tilde{W}_t denotes a standard Wiener process under Q . Following from the results in Duffie and Kan (1996), the time t price of a zero-coupon bond that matures at $t + \tau$ is exponential affine in the state-variables:

$$B(X_t, \tau) = e^{-A_1(\tau) - A_2(\tau)X_t}. \quad (61)$$

And, the corresponding zero-coupon bond yield is thus affine in the state-variables:

$$y(X_t, \tau) = \frac{A_1(\tau)}{\tau} + \frac{A_2(\tau)}{\tau} X_t \quad (62)$$

where $A_1(\tau)$ and $A_2(\tau)$ satisfy the ordinary differential equations

$$\frac{dA_2(\tau)}{d\tau} = -\tilde{K}' A_2(\tau) + \delta_X \quad (63)$$

$$\frac{dA_1(\tau)}{d\tau} = \tilde{\theta}' \tilde{K}' A_2(\tau) - \frac{1}{2} A_2(\tau)' \sigma_X \sigma_X A_2(\tau) + \delta_0 \quad (64)$$

with terminal conditions $A_2(\tau) = 0$ and $A_1(\tau) = 0$. Closed form expressions for $A_1(\tau)$ and $A_2(\tau)$ are available but are fairly messy and hence not reported here.

Campbell and Shiller (1988a) derive the following approximate expression for the log dividend-price ratio:

$$d_t - p_t = -\frac{k}{1 - \rho} + E_t \left[\sum_{j=0}^{\infty} \rho^j (-\Delta d_{t+j+1} + \hat{r}_{t+j+1} + \hat{\eta}_{t+j+1}) \right], \quad (65)$$

where Δd_{t+j+1} , \hat{r}_{t+j+1} and $\hat{\eta}_{t+j+1}$ denote log dividend growth, log interest rate and log excess return, respectively, from $t + j$ to $t + j + 1$. Furthermore, $\rho = 1/(1 + \exp(\overline{d - p}))$, with $\overline{d - p}$ denoting the average log dividend-price ratio, and $k = -\log(\rho) - (1 - \rho)\log(1 - \rho - 1)$. We simplify the analysis by assuming that expected dividend growth is constant and, hence, that all variation in the dividend-price ratio is due to variations in expected future interest rates and risk premia.

Based on (35), (36), (39) and (40) we can derive the expectations at time $t + j$ of \widehat{r}_{t+j+1} and $\widehat{\eta}_{t+j+1}$ as

$$E_{t+j}[\widehat{r}_{t+j+1}] = \int_{t+j}^{t+j+1} r(X_u) du = \delta_0 + \delta_X \nu X_{t+j} \quad (66)$$

$$E_{t+j}[\widehat{\eta}_{t+j+1}] = \int_{t+j}^{t+j+1} \left(-\frac{1}{2} \sigma_S^2 + \eta(X_u) \right) du = -\frac{1}{2} \sigma_S^2 + \eta_0 + \eta_X \nu X_{t+j} \quad (67)$$

where $\nu = \int_0^1 e^{-Ku} du$; see, e.g., Karatzas and Shreve (pp. 354-357, 1991). If we let $\Phi (= e^{-K})$ denote the discrete-time AR coefficient of X_t and g the mean log dividend growth, we have

$$d_t - p_t = -\frac{k}{1-\rho} + E_t \left[\sum_{j=0}^{\infty} \rho^j (-E_{t+j}[\Delta d_{t+j+1}] + E_{t+j}[\widehat{r}_{t+j+1}] + E_{t+j}[\widehat{\eta}_{t+j+1}]) \right] \quad (68)$$

$$= -\frac{k}{1-\rho} + E_t \left[\sum_{j=0}^{\infty} \rho^j \left(-g + \delta_0 + \delta_X \nu X_{t+1+j} - \frac{1}{2} \sigma_S^2 + \eta_0 + \eta_X \nu X_{t+1+j} \right) \right] \quad (69)$$

$$= -\frac{k + g + \frac{1}{2} \sigma_S^2 - \delta_0 - \eta_0}{1-\rho} + (\delta_X + \eta_X) \nu \left[\sum_{j=0}^{\infty} \rho^j \Phi^j X_t \right] \quad (70)$$

$$= -\frac{k + g + \frac{1}{2} \sigma_S^2 - \delta_0 - \eta_0}{1-\rho} + (\delta_X + \eta_X) \nu (I - \rho \Phi)^{-1} X_t. \quad (71)$$

	$\gamma = 2$			$\gamma = 5$			$\gamma = 10$		
	$T = 2$	$T = 10$	$T = 20$	$T = 2$	$T = 10$	$T = 20$	$T = 2$	$T = 10$	$T = 20$
α	1.59	1.48	1.36	0.63	0.56	0.50	0.31	0.28	0.24
α^*	1.62	1.62	1.62	0.65	0.65	0.65	0.32	0.32	0.32
Δ	-0.03	-0.14	-0.26	-0.02	-0.09	-0.15	-0.01	-0.04	-0.08
$\tilde{\alpha}^*$	1.62	1.62	1.62	0.65	0.65	0.65	0.32	0.32	0.32

Notes: Fraction of wealth allocated to stocks for varying risk-aversion coefficients, investment horizons and investor types. α denotes the allocation of a long-term investor who takes account of uncertainty about μ_t and learning. α^* denotes the allocation of a myopic investor who takes into account uncertainty about μ_t . Δ is the hedge demand of the long term investor. $\tilde{\alpha}^*$ denotes the allocation of a myopic investor who ignores uncertainty about μ_t .

Table 1: Allocations to stocks with and without learning

Discrete-time parameters		Continuous-time parameters	
r	0.0043	r	0.0516
$\bar{\mu}$	0.0035	θ	0.0531
ϕ_μ	0.9868	κ	0.1592
σ_x	0.0416	σ_S	0.1451
σ_μ	5.28 E-4	σ_η	0.0222
ρ	-0.9184	ρ	-0.9194

Notes: To the left are parameter estimates of the discrete-time return dynamics (32)-(34). These are in monthly terms. To the right are parameters of the continuous-time return dynamics (50)-(52). These were obtained from the discrete-time estimates and are in annual terms.

Table 2: Parameters of discrete-time and continuous-time models of section 4.2.

	$\gamma = 2$			$\gamma = 5$			$\gamma = 10$		
	$T = 2$	$T = 10$	$T = 20$	$T = 2$	$T = 10$	$T = 20$	$T = 2$	$T = 10$	$T = 20$
α	1.38	1.67	1.81	0.58	0.83	1.00	0.30	0.45	0.58
α^*	1.27	1.27	1.27	0.51	0.51	0.51	0.25	0.25	0.25
Δ	0.11	0.40	0.54	0.07	0.32	0.49	0.05	0.20	0.33
$\tilde{\alpha}$	1.44	1.92	2.16	0.62	1.08	1.41	0.32	0.62	0.89
$\tilde{\alpha}^*$	1.27	1.27	1.27	0.51	0.51	0.51	0.25	0.25	0.25
$\tilde{\Delta}$	0.17	0.65	0.89	0.11	0.57	0.90	0.07	0.37	0.64

Notes: Fraction of wealth allocated to stocks for varying risk-aversion coefficients, investment horizons and investor types when the risk-premium equals its steady state mean. α and α^* denote allocations for a long term and a myopic investor, respectively, who filter out the risk premium by observing realized stock returns. Δ is the hedge demand of the long term investor. $\tilde{\alpha}$ and $\tilde{\alpha}^*$ denote allocations for a long term and myopic investor, respectively, who observe the true risk premium. $\tilde{\Delta}$ is the hedge demand of this long term investor.

Table 3: Allocations to stocks with and without observable risk premium

	annually	quarterly	monthly	weekly	daily	continuous
α	0.8486	0.8343	0.8309	0.8296	0.8292	0.8292
α^*	0.5434	0.5139	0.5074	0.5049	0.5043	0.5041
Δ	0.3052	0.3204	0.3235	0.3247	0.3249	0.3251
$\tilde{\alpha}$	1.1063	1.0860	1.0810	1.0791	1.0786	1.0785
$\tilde{\alpha}^*$	0.5627	0.5185	0.5089	0.5052	0.5043	0.5041
$\tilde{\Delta}$	0.5436	0.5675	0.5721	0.5739	0.5743	0.5744

Notes: Fraction of wealth allocated to stocks for varying rebalancing intervals and investor types when the risk-premium equals its steady state mean. α and α^* denote allocations for a long term and a myopic investor, respectively, who filter out the risk premium by observing realized stock returns. Δ is the hedge demand of the long term investor. $\tilde{\alpha}$ and $\tilde{\alpha}^*$ denote allocations for a long term and myopic investor, respectively, who observe the true risk premium. $\tilde{\Delta}$ is the hedge demand of this long term investor. Investment horizon is 10 years and $\gamma = 5$.

Table 4: Allocations to stocks for various rebalancing intervals

	constant	X_1	X_2	X_3
δ	0.0396	0.0178	0.0056	0
	<i>0.0209</i>	<i>0.0017</i>	<i>0.0053</i>	-
λ_1	-0.4492	0.1344	0.3493	0
	<i>0.1752</i>	<i>0.1101</i>	<i>0.0155</i>	-
λ_2	-0.3076	-0.2286	-0.1925	0
	<i>0.1875</i>	<i>0.1205</i>	<i>0.1143</i>	-
η	0.0484	-0.0150	-0.0114	-0.0092
	<i>0.0180</i>	<i>0.0126</i>	<i>0.0147</i>	<i>0.0072</i>
ζ	0.0297	0.0146	0.0074	0
	<i>0.0142</i>	<i>0.0024</i>	<i>0.0045</i>	-
K		0.1399	0	0
		<i>0.0835</i>	-	-
		0.4209	0.4941	0
		<i>0.1269</i>	<i>0.1897</i>	-
		0	0	0.0344
		-	-	<i>0.0437</i>
σ	0.0365			
	<i>0.0005</i>			

Notes: Maximum-likelihood estimates. White (1980) asymptotic heteroscedasticity-consistent standard errors in italics. Estimates are reported in pro-anno terms. σ denotes the standard deviation on the pricing error in the bond yield and dividend-price ratio relations.

Table 5: Parameter estimates

	dZ_1	dZ_2	dZ_3	dZ_4	dZ_5
$\sigma_{X,1}$	1	0	0	0	0
	-	-	-	-	-
$\sigma_{X,2}$	0	1	0	0	0
	-	-	-	-	-
$\sigma_{X,3}$	0	0	1	0	0
	-	-	-	-	-
$\sigma_{S,1}$	-0.0219	0.0130	0.1396	0.0306	0
	<i>0.0075</i>	<i>0.0093</i>	<i>0.0077</i>	<i>0.0320</i>	-
$\sigma_{\Pi,1}$	0.0007	-0.0009	-0.0009	0.0004	0.0086
	<i>0.0006</i>	<i>0.0005</i>	<i>0.0010</i>	<i>0.0041</i>	<i>0.0004</i>

Notes: Maximum-likelihood estimates. Asymptotic heteroscedasticity-consistent standard errors in italics. Estimates are reported in pro-anno terms.

Table 6: Parameter estimates in volatility matrix

	Sample			Implied		
	mean	std. dev.	S.R.	mean	std. dev.	S.R.
Real rate	0.0142	0.0101		0.0103	0.0097	
1yr bond	0.0042	0.0155	0.2700	0.0080	0.0154	0.5209
5yr bond	0.0104	0.0605	0.1710	0.0186	0.0551	0.3378
10yr bond	0.0129	0.1001	0.1290	0.0210	0.0988	0.2130
Stock	0.0644	0.1192	0.5405	0.0416	0.1198	0.3475

Notes: Unconditional moments of the real interest rate, the 1yr, 5yr and 10yr excess bond returns and the excess stock return. S.R. denotes Sharpe Ratio.

Table 7: Unconditional moments

	state 1	state 2	state 3
real rate	0.991	-0.037	0.000
1yr riskp	-0.000	-1.000	-0.000
5yr riskp	-0.000	-1.000	-0.000
10yr riskp	-0.000	-1.000	-0.000
stock riskp	-0.619	-0.483	-0.619
1yr return	-0.911	-0.412	-0.000
5yr return	-0.999	0.047	-0.000
10yr return	-0.959	0.283	-0.000
stock return	-0.029	0.031	0.987

Notes: Instantaneous correlations between innovations to state-variables and innovations to the real rate, bond and stock risk premia and realized bond and stock returns

Table 8: Correlations among innovations to central variables

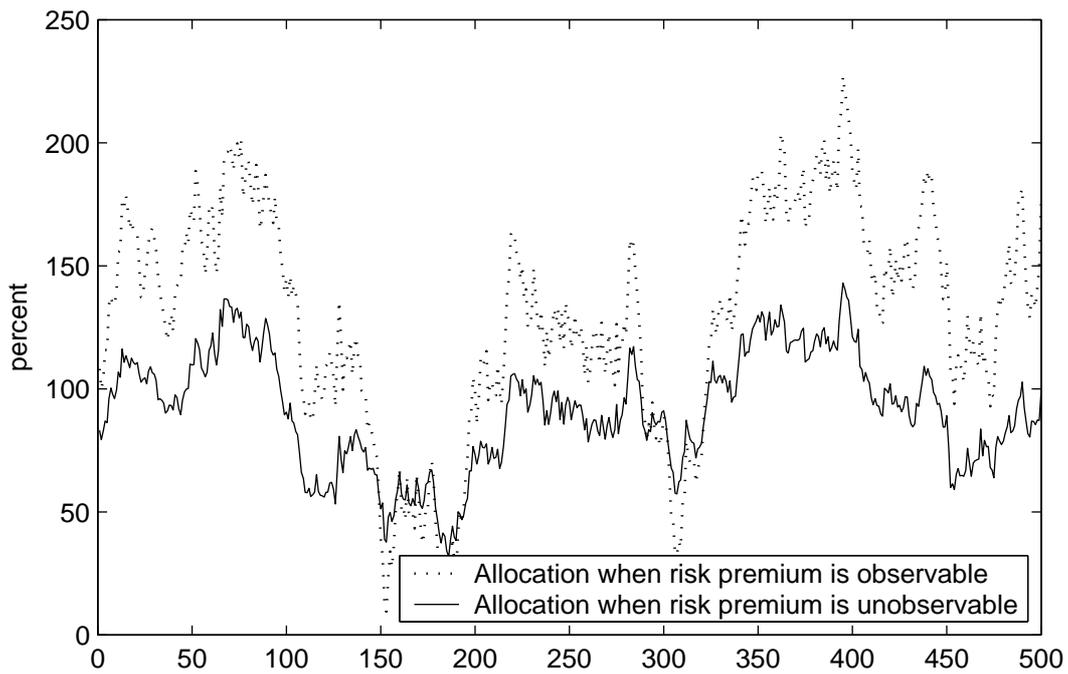


Figure 1: Allocations to stocks with and without observable risk premium

The figure shows the fraction of wealth allocated to stocks when the risk premium is observable, respectively, unobservable. In the latter case the allocation is based on a Kalman filtered estimate of the risk premium. The investment horizon is 10 years and $\gamma = 5$.

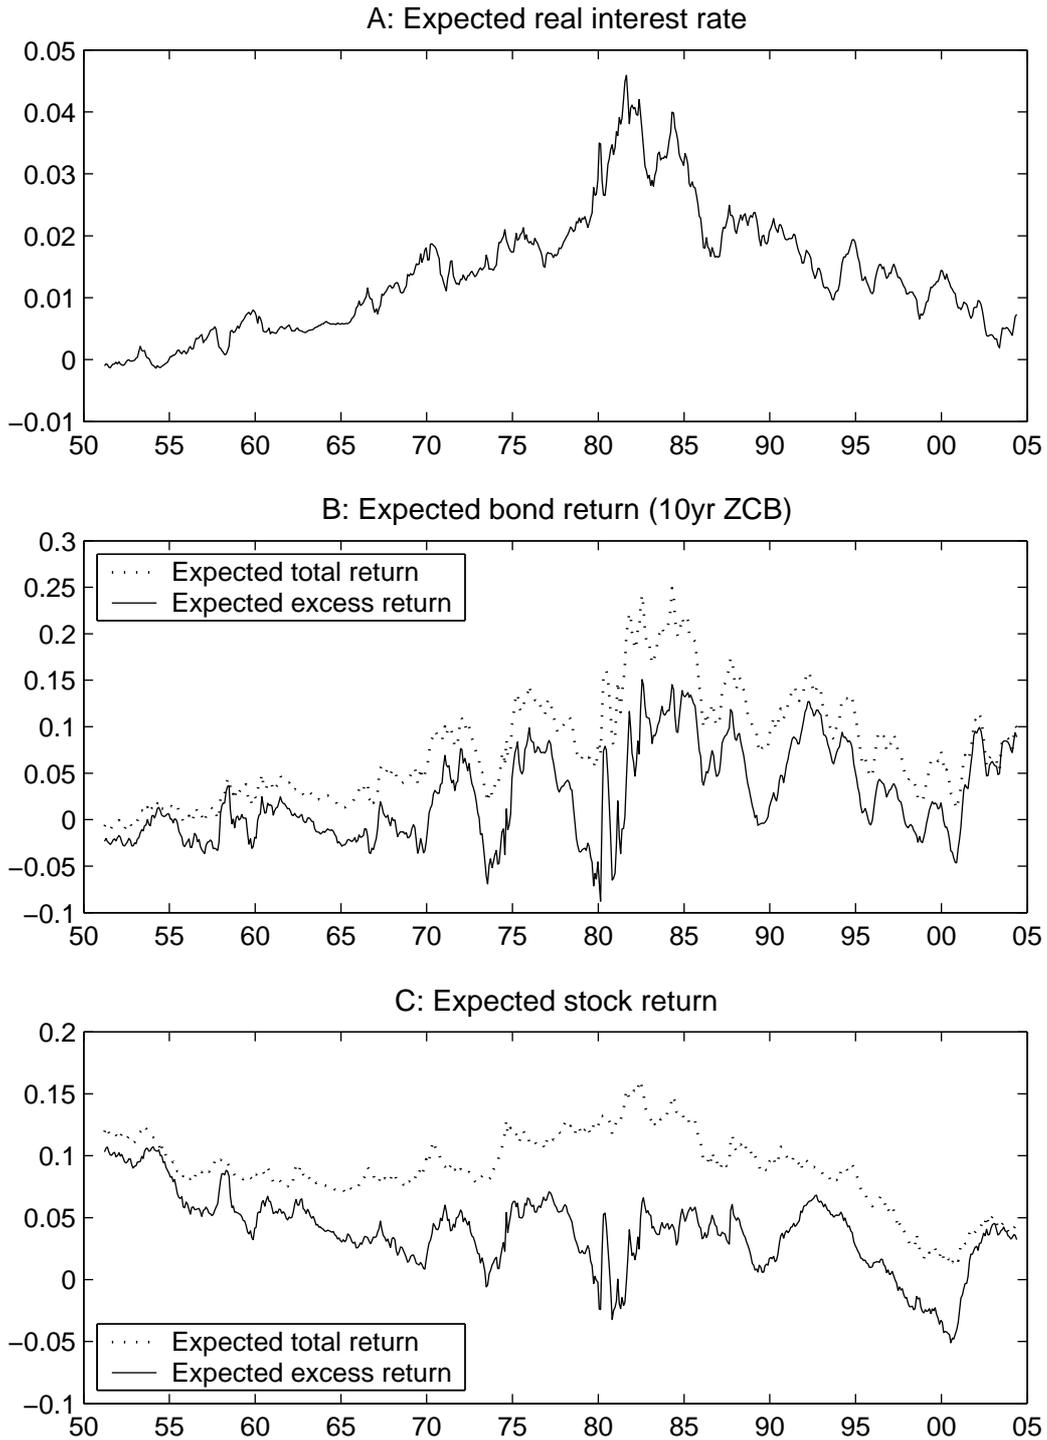


Figure 2: Expected real interest rate and excess bond and stock return

Panel A shows fitted values for the instantaneous expected real interest rate. Panel B shows fitted values for the instantaneous expected total and excess returns on 10yr bonds. Panel C shows fitted values for the instantaneous expected total and excess returns on stocks. The fitted values are calculated on the basis of the Kalman filtered smoothed state-variables. Units are annual.

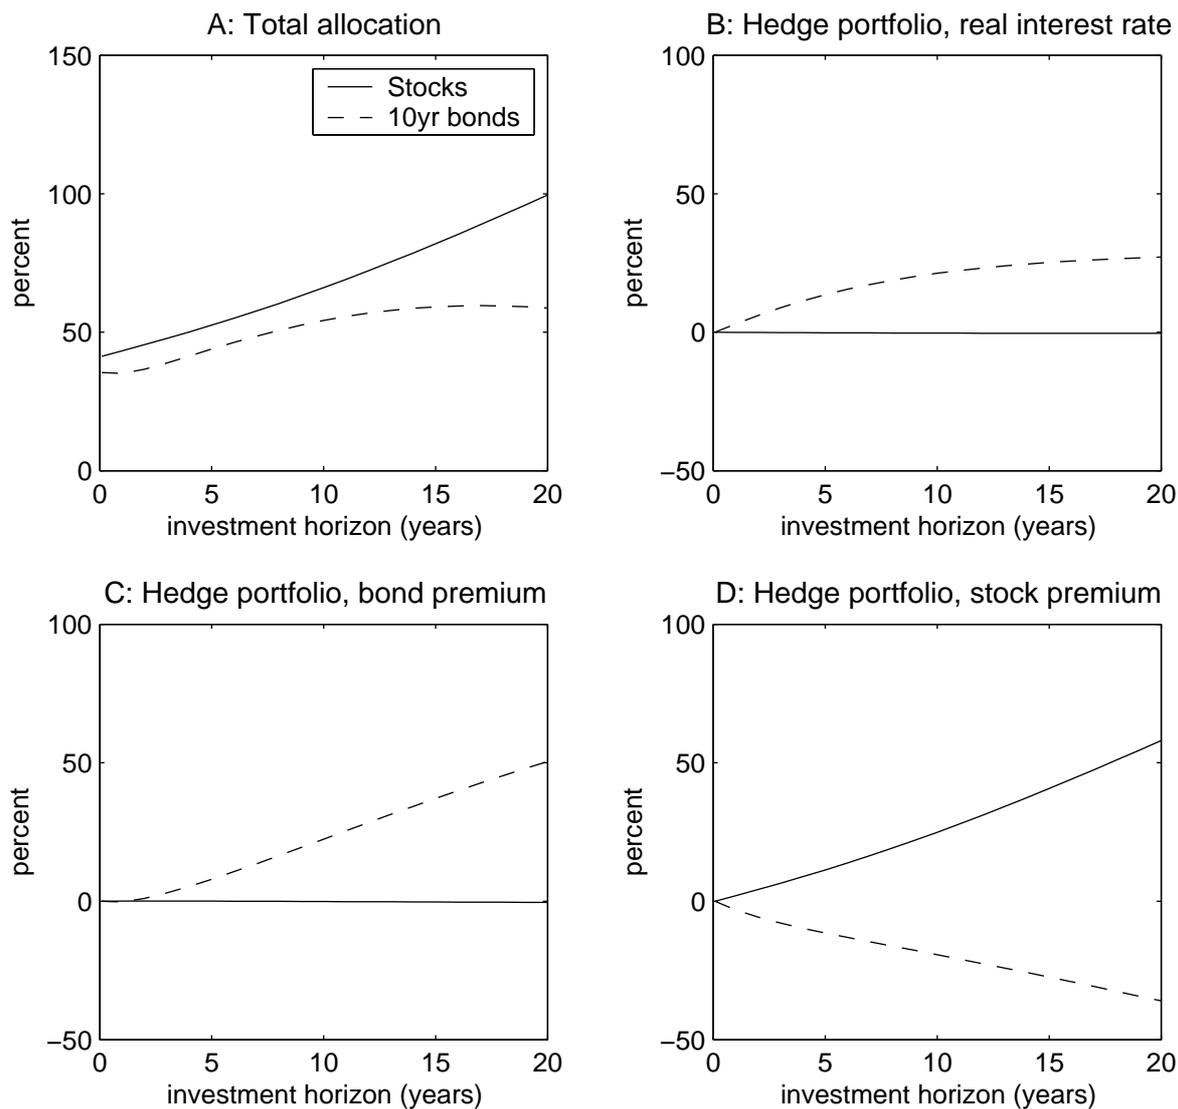


Figure 3: Asset allocation as function of investment horizon

Optimal asset allocation as a function of the investment horizon when the investor is assumed to invest in stocks, 10yr zero-coupon bonds and cash. We assume $\gamma = 5$ and the state vector equals its unconditional mean. Panel A shows the total allocation to stocks and bonds. The horizon effect is due to long-term investors hedging variations in real interest rates and risk premia on stocks and bonds. Panel B shows the hedge portfolio when the investor only hedges variations in the real interest rate and treats risk premia on stocks and bonds as constants and equal to their unconditional means. Panel C shows the hedge portfolio when the investor only hedges variations in the bond risk premium and treats the expected real interest rate and the risk premium on stocks as constants and equal to their unconditional means. Finally, Panel D shows the hedge portfolio when the investor only hedges variations in the stock risk premium and treats the expected real interest rate and the risk premium on bonds as constants and equal to their unconditional means. Throughout, solid lines denotes stock allocations and dashed lines denotes bond allocations

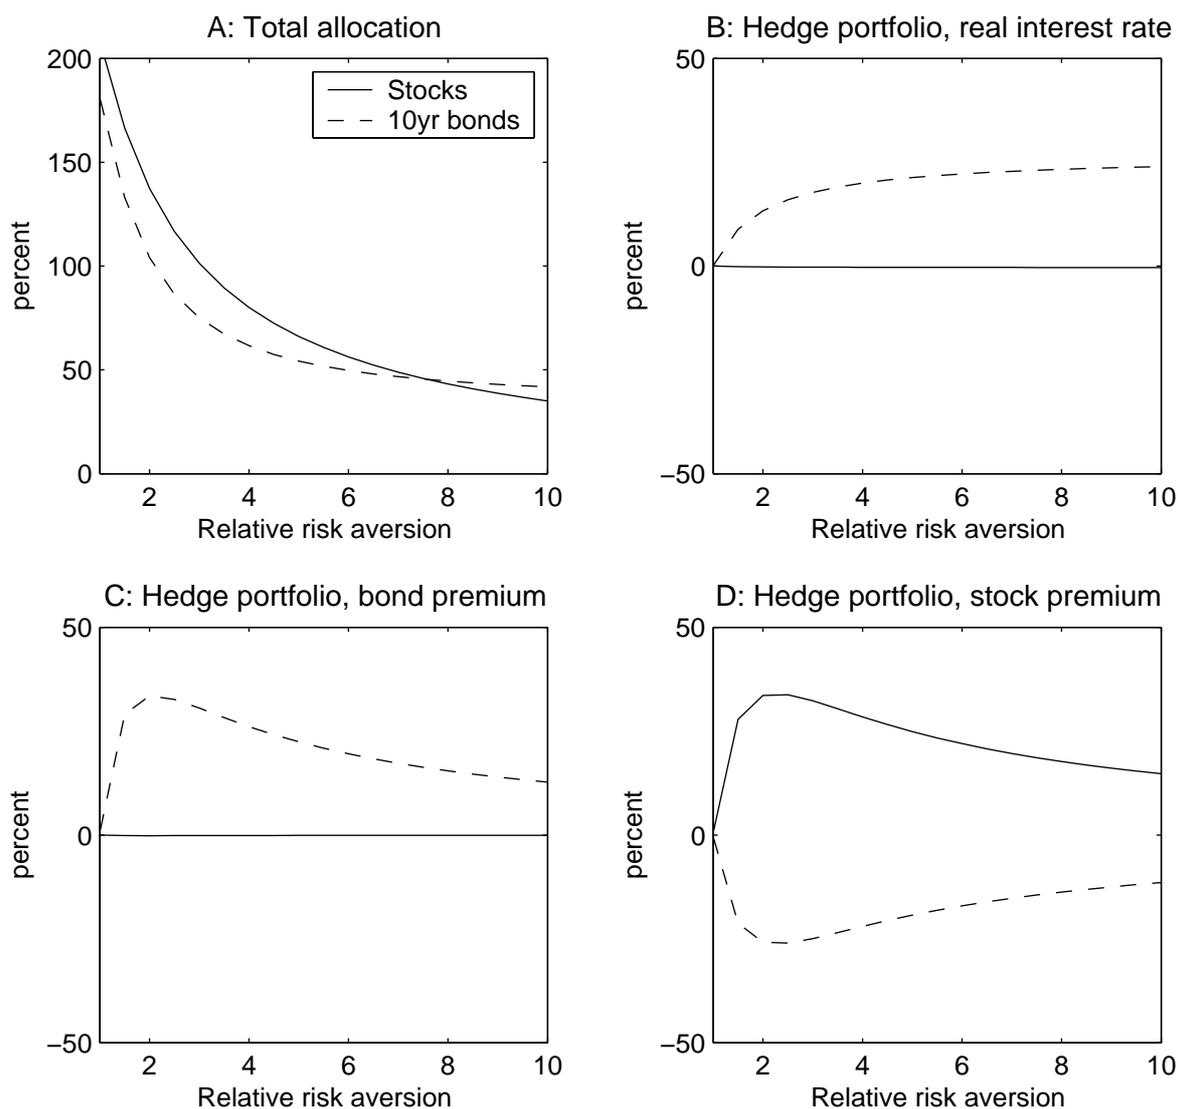


Figure 4: Asset allocation as function of relative risk aversion

Optimal asset allocation as a function of the degree of relative risk aversion when the investor is assumed to invest in stocks, 10yr zero-coupon bonds and cash. We assume an investment horizon of 10 years and the state vector equals its unconditional mean. Panel A shows the total allocation to stocks and bonds. Panel B-D focus on the hedge portfolios. Panel B shows the hedge portfolio when the investor only hedges variations in the real interest rate and treats risk premia on stocks and bonds as constants and equal to their unconditional means. Panel C shows the hedge portfolio when the investor only hedges variations in the bond risk premium and treats the expected real interest rate and the risk premium on stocks as constants and equal to their unconditional means. Finally, Panel D shows the hedge portfolio when the investor only hedges variations in the stock risk premium and treats the expected real interest rate and the risk premium on bonds as constants and equal to their unconditional means. Throughout, solid lines denotes stock allocations and dashed lines denotes bond allocations.

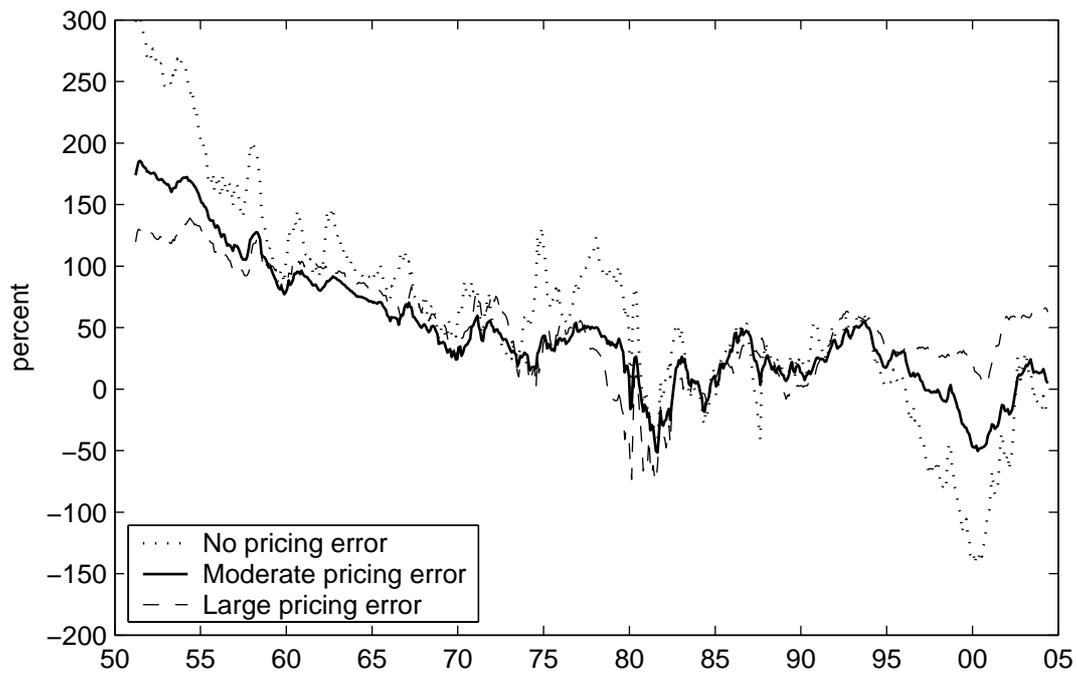


Figure 5: Allocations to stocks for various degrees of confidence in the dividend-yield relation. Allocations to stocks over time for different assumptions about pricing errors on the dividend-yield relation in the measurement equation. The dotted line shows the allocations when the pricing error is assumed to be zero. The solid line shows the allocations in the base-line case where the percentage standard deviation of the pricing errors is equal across the bond yield and dividend yield relations. The dashed line shows the allocations when the pricing error on the dividend yield relation is increased by a factor two relative to the base-line case. We assume an investment horizon of ten years and $\gamma = 5$.

References

- Barberis, N. (2000). “Investing for the Long Run when Returns are Predictable”. *Journal of Finance*, 55(1), 225–264.
- Bawa, V. S.; Brown, S., & Klein, R. W. (1979). *Estimation Risk and Optimal Portfolio Choice*. North-Holland, Amsterdam.
- Brennan, M. J. (1998). “The Role of Learning in Dynamic Portfolio Decisions”. *European Finance Review*, 1, 295–306.
- Brennan, M. J.; Schwartz, E. S., & Lagnado, R. (1997). “Strategic Asset Allocation”. *Journal of Economic Dynamics and Control*, 21, 1377–1403.
- Brennan, M. J. & Xia, Y. (2000). “Stochastic Interest Rates and Bond-Stock Mix”. *European Finance Review*, 4, 197–210.
- Brennan, M. J. & Xia, Y. (2002). “Dynamic Asset Allocation under Inflation”. *Journal of Finance*, 57(3), 1201–1238.
- Brennan, M. J. & Xia, Y. (2004). “Persistence, Predictability, and Portfolio Planning”. Working paper, Anderson Graduate School of Management, UCLA.
- Campbell, J. Y.; Chacko, G.; Rodriguez, J., & Viceira, L. M. (2004). “Strategic asset allocation in a continuous-time VAR model”. *Journal of Economic Dynamics and Control*, 28, 2195–2214.
- Campbell, J. Y.; Chan, Y. L., & Viceira, L. M. (2002). “Appendix to a multivariate model of strategic asset allocation”. Unpublished manuscript, Harvard University.
- Campbell, J. Y.; Chan, Y. L., & Viceira, L. M. (2003). “A multivariate model of strategic asset allocation”. *Journal of Financial Economics*, 67, 41–80.
- Campbell, J. Y. & Shiller, R. J. (1988a). “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors”. *Review of Financial Studies*, 1(3), 195–228.
- Campbell, J. Y. & Shiller, R. J. (1988b). “Stock Prices, Earnings, and Expected Dividends”. *Journal of Finance*, 43(3), 661–676.

- Campbell, J. Y. & Viceira, L. M. (1999). “Consumption and Portfolio Decisions when Expected Returns are Time-Varying”. *The Quarterly Journal of Economics*, 114, 433–495.
- Campbell, J. Y. & Viceira, L. M. (2001). “Who Should Buy Long-Term Bonds?”. *American Economic Review*, 91(1), 99–127.
- Campbell, J. Y. & Viceira, L. M. (2002). *Strategic Asset Allocation*. Oxford University Press.
- Chacko, G. & Viceira, L. M. (1999). “Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets”. Working paper, Harvard University.
- Dai, Q. & Singleton, K. J. (2000). “Specification analysis of affine term structure models”. *Journal of Finance*, 55, 1943–1978.
- Dai, Q. & Singleton, K. J. (2002). “Expectations puzzles, time-varying risk premia, and affine models of the term structure”. *Journal of Financial Economics*, 63, 415–442.
- Detemple, J. (1986). “Asset pricing in an economy with incomplete information”. *Journal of Finance*, 41, 383–392.
- Dothan, M. D. & Feldman, D. (1986). “Equilibrium Interest Rates and Multiperiod Bonds in a Partially Observable Economy”. *Journal of Finance*, 41(2), 369–382.
- Duffee, G. R. (2002). “Term premia and interest rate forecasts in affine models”. *Journal of Finance*, 57, 369–443.
- Duffie, J. D. & Kan, R. (1996). “A Yield-Factor Model of Interest Rates”. *Mathematical Finance*, 6, 379–406.
- Gennotte, G. (1986). “Optimal Portfolio Choice Under Incomplete Information”. *Journal of Finance*, 41, 733–749.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press, Princeton, New Jersey.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.
- Karatzas, I. & Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus*. New York: Springer-Verlag, 2 edition.

- Kim, T. S. & Omberg, E. (1996). “Dynamic Nonmyopic Portfolio Behavior”. *Review of Financial Studies*, 9(1), 141–161.
- Klein, R. W. & Bawa, V. S. (1976). “The Effect of Estimation Risk on Optimal Portfolio Choice”. *Journal of Financial Economics*, 3, 215–231.
- Lipster, R. S. & Shiryaev, A. N. (2001). *Statistics of Random Processes I & II*. Springer-Verlag, Berlin.
- Liu, J. (1999). “Portfolio Selection in Stochastic Environments”. Working paper, Stanford University.
- Markowitz, H. (1952). “Portfolio Selection”. *Journal of Finance*, 7, 77–91.
- McCulloch, J. H. (1990). “U.S. Term Structure Data, 1946-87”. *Handbook of Monetary Economics*, Volume I, 672–715.
- McCulloch, J. H. & Kwon, H.-C. (1993). “U.S. Term Structure Data, 1947-91”. Working paper, Economics Department, Ohio State University.
- Merton, R. C. (1969). “Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case”. *Review of Economics and Statistics*, 51, 247–257.
- Merton, R. C. (1971). “Optimum Consumption and Portfolio Rules in a Continuous-Time Model”. *Journal of Economic Theory*, 3, 373–413.
- Merton, R. C. (1973). “An Intertemporal Capital Asset Pricing Model”. *Econometrica*, 41, 867–887.
- Pennacchi, G. G. (1991). “Identifying the Dynamics of Real Interest Rates and Inflation: Evidence Using Survey Data”. *Review of Financial Studies*, 4(1), 53–86.
- Rodriguez, J. F. (2002). “Hedging Demands, Incomplete Markets, and Imperfect Information”. Working paper, Sloan School of Managements, Massachusetts Institute of Technology.
- Samuelson, P. A. (1969). “Lifetime Portfolio Selection by Dynamic Stochastic Programming”. *Review of Economics and Statistics*, 51, 239–246.
- Sangvinatsos, A. & Wachter, J. A. (2003). “Does the Failure of the Expectations Hypothesis Matter for Long-Term Investors?”. Working paper, Stern School of Business, New York University.

- Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press.
- Sørensen, C. (1999). “Dynamic Asset Allocation and Fixed Income Management”. *Journal of Financial and Quantitative Analysis*, 34(4), 513–531.
- Wachter, J. A. (2002). “Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets”. *Journal of Financial and Quantitative Analysis*, 37(1), 63–91.
- White, H. (1980). “A heteroscedasticity-consistent covariance estimator and a direct test for heteroscedasticity”. *Econometrica*, 48, 817–838.
- Williams, J. T. (1977). “Capital asset prices with heterogenous beliefs”. *Journal of Financial Economics*, 5, 219–240.

- WP 2000-1 Bjarne Astrup Jensen and Carsten Sørensen: Paying for minimum interest rate guarantees: Who should compensate who?
- WP 2000-2 Jan Jakobsen and Ole Sørensen: Decomposing and testing Long-run Returns with an application to initial public offerings in Denmark.
- WP 2000-3 Jan Jakobsen and Torben Voetmann: Volatility-Adjusted Performance An Alternative Approach to Interpret Long-Run Returns.
- WP 2000-4 Jan Jakobsen and Torben Voetmann: Post-Acquisition Performance in the Short and Long-Run Evidence from the Copenhagen Stock Exchange 1993-1997.
- WP 2000-5 Ken L. Bechmann and Johannes Raaballe: A Regulation of Bids for Dual Class Shares. Implication: Two Shares – One Price.
- WP 2000-6 Torben Voetmann: Changes in the Bid-Ask Components Around Earnings Announcements: Evidence from the Copenhagen Stock Exchange.
- WP 2000-7 Henrik Lando: The Optimal Standard of Proof in Criminal Law When Both Fairness and Deterrence Are Social Aims.
- WP 2000-8 Jesper Rangvid and Carsten Sørensen: Convergence in the ERM and Declining Numbers of Common Stochastic Trends.
- WP 2000-9 Claus Munk and Carsten Sørensen: Optimal Consumption and Investment Strategies with Stochastic Interest Rates.
- WP 2000-10 Henrik Lando and Caspar Rose: On Specific Performance in Civil Law and Enforcement Costs.
- WP 2000-11 Henrik Lando: Ny lov om jordforurening i økonomisk belysning.
-
- WP 2001-1 Michael Møller, Claus Parum og Thomas Sørensen: Den ny pensionsafkastbeskatningslov.
- WP 2001-2 Bjarne Astrup Jensen: Mean variance efficient portfolios by linear programming: A review of some portfolio selection criteria of Elton, Gruber and Padberg.
- WP 2001-3 Caspar Rose: Impact of Investor Meetings/Presentations on Share

Prices, Insider Trading and Securities Regulation.

- WP 2001-4 Caspar Rose: Corporate Financial Performance and the Use of Takeover Defenses.
- WP 2001-5 Shubhashis Gangopadhyay and Clas Wihlborg: The Impact of Bankruptcy Rules on Risky Project Choice and Skill Formation under Credit Rationing.
- WP 2001-6 Claus Munk, Carsten Sørensen & Tina Nygaard Vinther: Portfolio Choice under Inflation: Are Popular Recommendations Consistent with Rational Behaviour?
- WP 2001-7 Ken L. Bechmann: Evidence on the Limits of Arbitrage: Short Sales, Price Pressure, and the Stock Price Response to Convertible Bond Calls.
- WP 2001-8 Michael Møller & Caspar Rose: Legal pre-emption rights as call-options, redistribution and efficiency loss.
- WP 2001-9 Peter Raahauge: Empirical Rationality in the Stock Market.

- WP 2002-1 Bjarne Astrup Jensen: On valuation before and after tax in no arbitrage models: Tax neutrality in the discrete time model.
- WP 2002-2 Ken L. Bechmann: Price and Volume Effects Associated with Changes in the Danish Blue-Chip Index – The KFX Index.
- WP 2002-3 Steen Thomsen and Caspar Rose: Foundation ownership and financial performance. Do companies need owners?
- WP 2002-4 Martin Richter and Carsten Sørensen: Stochastic Volatility and Seasonality in Commodity Futures and Options: The Case of Soybeans.
- WP 2002-5 Caspar Rose: Aktiemarkedets reaktion på indførelsen af incitamentsprogrammer.
- WP 2002-6 Caspar Rose: Impact of Takeover Defenses on Managerial Incentives.
- WP 2002-7 Ken L. Bechmann og Peter Løchte Jørgensen: Options aflønning i danske børsnoterede selskaber.
- WP 2002-8 Jesper Rangvid: Output and Expected Returns – a multicountry study.
- WP 2002-9 Jonas Aziz Bhatti og Michael Møller: Pensionsafkastbeskatning og optimal porteføljesammensætning.

-
- WP 2003-1 Bjarne Florentsen, Michael Møller and Niels Chr. Nielsen: Reimbursement of VAT on written-off Receivables.
- WP 2003-2: Ken L. Bechmann and Peter Løchte Jørgensen: The Value and Incentives of Option-based Compensation in Danish Listed Companies.
-
- WP 2004-1 Ken L. Bechmann and Johannes Raaballe: The Difference Between Stock Splits and Stock Dividends – Evidence from Denmark.
- WP 2004-2 Caspar Rose: Bestyrelsessammensætning og finansiel performance i danske børsnoterede virksomheder - Er Nørbyrapportens anbefalinger til gavn for aktionærene?
- WP 2004-3 Jens Lunde: Lack of balance in after-tax returns - lack of tenure neutrality. The Danish case.
- WP 2004-4 Peter Raahauge: Upper Bounds on Numerical Approximation Errors.
- WP 2004-5 Peter Raahauge: Higher-Order Finite Element Solutions of Option Prices.
- WP 2004-6 Anders Holst & Morten Nalholm: Modelling Callable Annuity Bonds with Interest-Only Optionality.
- WP 2004-7 Bent Jesper Christensen & Peter Raahauge: Latent Utility Shocks in a Structural Empirical Asset Pricing Model.
- WP 2004-8 Carsten Sørensen and Anders Bjerre Trolle: Dynamic asset allocation and latent variables.