

Strategic asset allocation for long-term investors: Parameter uncertainty and prior information

Citation for published version (APA):

Hoevenaars, R. P. P. M., Molenaar, R. D. J., Schotman, P. C., & Steenkamp, T. B. M. (2014). Strategic asset allocation for long-term investors: Parameter uncertainty and prior information. *Journal of Applied Econometrics*, 29(3), 353-376. <https://doi.org/10.1002/jae.2331>

Document status and date:

Published: 01/01/2014

DOI:

[10.1002/jae.2331](https://doi.org/10.1002/jae.2331)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

STRATEGIC ASSET ALLOCATION FOR LONG-TERM INVESTORS: PARAMETER UNCERTAINTY AND PRIOR INFORMATION

ROY P. P. M. HOEVENAARS^a, RODERICK D. J. MOLENAAR^b, PETER C. SCHOTMAN^{c*}
AND TOM B. M. STEENKAMP^d

^a *Blenheim Capital Management BV, Amsterdam, The Netherlands*

^b *Robeco Investments, Rotterdam, The Netherlands*

^c *Maastricht University, and Netspar, Maastricht, The Netherlands*

^d *Robeco Investments, and Free University Amsterdam, The Netherlands*

SUMMARY

We study the effect of parameter uncertainty on the long-run risk for three asset classes: stocks, bills and bonds. Using a Bayesian vector autoregression with an uninformative prior we find that parameter uncertainty raises the annualized long-run volatilities of all three asset classes proportionally with the same factor relative to volatilities that are conditional on maximum likelihood parameter estimates. As a result, the horizon effect in optimal asset allocations is much weaker compared to models in which only equity returns are subject to parameter uncertainty. Results are sensitive to alternative informative priors, but generally the term structure of risk for stocks and bonds is relatively flat for investment horizons up to 15 years. Copyright © 2013 John Wiley & Sons, Ltd.

Received 04 April 2011; Revised 27 December 2012

1. INTRODUCTION

Should long-term investors hold a different portfolio than short-term investors? A large part of the answer depends on the risk characteristics of alternative asset classes at different investment horizons. For this purpose Campbell and Viceira (2005) estimate a vector autoregression (VAR) to model the return dynamics of stocks, bonds and bills. For equity they find mean reversion, which leads to a decrease of the per-period variance of stocks from 15% for a 1-year horizon to 8% for a 30-year investment horizon. Real returns on short-term T-bills, however, show increasing risk at longer investment horizons. As a result, Campbell *et al.* (2003), among others, conclude that long-term investors should hold more equity.

These risk estimates are, however, conditional on the parameter estimates of the VAR model and subject to substantial parameter uncertainty. For equity returns various Bayesian studies, e.g. Barberis (2000) and Pastor and Stambaugh (2012), show that accounting for parameter uncertainty raises the long-run risk of equity and this may be quantitatively important. They explain that uncertainty about the expected return is particularly important for long-run investors, since the estimation error is perfectly correlated over time and therefore quickly increases with the investment horizon. Pastor and Stambaugh (2012) even conclude that the effect of mean reversion in stock returns is more than offset by uncertainties about the expected stock return.

The implications for portfolio choice are not clear, however. Both studies, as well as Avramov (2002) and Brandt *et al.* (2005), among others, assume that long-term investors can also invest in a real risk-free asset. When a real long-term bond is not available, a long-term investor's alternative to equity consists mainly of nominal bonds of various maturities. Since the real returns on nominal bonds are

* Correspondence to: Peter C. Schotman, Maastricht University, and Netspar, Maastricht, The Netherlands. E-mail: p.schotman@maastrichtuniversity.nl

subject to inflation and interest rate risk, parameter uncertainty will also increase their long-run risk. For portfolio choice we need to know the impact of parameter uncertainty on all variances and covariances.

We extend the Campbell and Viceira (2005) model by adding Bayesian parameter uncertainty for a range of priors. This generates a term structure of risk for all three asset classes and their correlations. This way we also extend the Barberis (2000) and Pastor and Stambaugh (2012) analyses, which focus entirely on equity risk. We concentrate on long horizons, up to 15 years, and this way we add to the Bayesian analysis of Wachter and Warusawitharana (2009), who show the good out-of-sample performance of Bayesian portfolio rules for a short investment horizon of three months. An important difference with many earlier studies is our focus on long-horizon returns and the riskiness of the benchmark asset. The long-term investor in our model cannot invest in a long-term real risk-free asset and thus faces uncertainty about the returns of all asset classes.

As priors we consider alternative views on the unconditional means of returns and state variables in the VAR. With an uninformative prior the long-term risk increases relative to estimates conditional on maximum likelihood estimates. For short-term investment horizons, the effect of parameter uncertainty is small. Long-run risk increases not only for equity but for all three asset classes, and the increase is almost proportionally. The combined effect of predictability in returns and parameter uncertainty leads to a nearly flat term structure of risk for stocks and bonds up to a horizon of 15 years. For bills the term structure of risk is upward sloping. For the correlations between the three asset classes there is no visible effect of parameter uncertainty.

Not surprisingly, more informative priors mostly imply smaller effects of parameter uncertainty. But informative priors that move expected returns away from the sample mean also increase the term structure of risk relative to maximum likelihood estimates of the VAR model of returns. For example, when the prior view is informative about a low long-run dividend yield, this also affects the vector autoregressive dynamics by shifting more probability mass to near unit root behaviour of the dividend yield, which in turn increases the long-term predictive variance of equity returns.

Since the term structure of risk is relatively flat under an uninformative prior, optimal Bayesian portfolios are similar for different investment horizons. Parameter uncertainty raises the overall risk and shifts the optimal Bayesian portfolios away from the most risky asset class (equity) towards the least risky asset class (bills). The differences across alternative priors can be large. As an extreme case we analyse a 'pessimistic' prior that imposes a low equity premium and a substantial bond premium. Obviously this prior drastically alters the optimal portfolio allocation for moderately risk-averse investors. Less obviously, it also affects the term structure of risk and thus even the minimum variance portfolio or an inflation hedge portfolio. This latter effect is due to the near unit roots in the VAR that are affected by the prior on the unconditional means.

Various other types of prior have been proposed in the literature. Most are related to the predictability of excess equity returns. These priors are informative about either the predictive R^2 in the return prediction regressions (e.g. Kandel and Stambaugh, 1996; Wachter and Warusawitharana, 2009), the covariance of predicted and unpredicted excess returns (Pastor and Stambaugh, 2009), or the Sharpe ratio (Shanken and Tamayo, 2004). Since these priors mostly conform with data evidence, they most likely reduce the effect of parameter uncertainty relative to our uninformative prior.

All our results are based on a first-order VAR with homoskedastic errors. Some have challenged the VAR(1) and proposed alternative models. Examples are time-varying parameters (e.g. Xia, 2001; Dangl and Halling, 2012; Johannes *et al.*, 2012), structural breaks and regimes shifts (e.g. Pesaran and Timmermann, 2002; Guidolin and Timmermann, 2007), alternative predictors (e.g. Cremers, 2002) and dynamic misspecification (Pastor and Stambaugh, 2012). All these priors deal with equity prediction models and are not directly applicable in a general VAR with multiple asset classes without a risk-free rate. We discuss these alternatives in a separate section after presenting our results.

The remainder of this paper is organized as follows. Sections 2 and 3 analyse the effect of Bayesian parameter uncertainty on measures of long-run risk and the composition of long-term portfolios. In Sections 4 and 5 we describe the details of the Bayesian estimation and simulation of the first-order VAR model, together with the data and the most salient estimation results. Results for long-term risk and portfolio allocation are presented in Sections 6 and 7. In Section 8 we discuss alternative priors and models. Finally, Section 9 concludes.

2. TERM STRUCTURE OF RISK

Let r_t be a vector of continuously compounded real returns, more specifically the (3×1) vector with real returns on a short term T-bill, long-term bonds and stocks, respectively:

$$r_t' = (r_{tb,t} \quad r_{b,t} \quad r_{s,t}) \quad (1)$$

For much of the analysis we will be interested in cumulative returns over k periods:

$$R_{t+k} = \sum_{j=1}^k r_{t+j} \quad (2)$$

Following Campbell *et al.* (2003) and Campbell and Viceira (2005), among others, we describe the return dynamics by a first-order vector autoregression (VAR) for returns and state variables:

$$y_t = \mu + \mathbf{B}(y_{t-1} - \mu) + \varepsilon_t \quad (3)$$

containing the variables

$$y_t = \begin{pmatrix} r_t \\ z_t \end{pmatrix}$$

with z_t a vector of observed state variables that help predict returns, and ε_t a vector of shocks with zero mean and covariance matrix Σ . The VAR has parameter vector $\theta' = (\mu' \text{vec}(\mathbf{B})' \text{vech}(\Sigma)')$ and implies conditional return distributions

$$p(R_{t+k}|y_t, \theta) \quad (4)$$

and more specifically $p(r_{t+1}|y_t, \theta)$ for $k=1$.

For inference we use Bayesian methods and define a prior $p(\theta)$. Adding historical sample data \mathbf{Y} we obtain the posterior $p(\theta|\mathbf{Y})$ and the predictive distributions

$$p(R_{t+k}|y_t, \mathbf{Y}) = \int p(R_{t+k}|y_t, \theta) p(\theta|\mathbf{Y}) d\theta \quad (5)$$

We use these quantities to estimate long-run risk and construct portfolios for long-run investors.

Our measure of long-run risk is the per period variance of returns as in Campbell and Viceira (2005). As a function of the investment horizon k it is called the term structure of risk, which for each asset class i is defined as

$$\sigma_i^2(k) = \frac{1}{k} \text{var}[R_{i,t+k}|y_t] \quad (6)$$

In Campbell and Viceira (2005) this variance is computed conditional on the estimated VAR parameters $\hat{\theta}$ using the conditional densities $p(R_{t+k}|y_t, \theta = \hat{\theta})$. In the Bayesian framework we account for parameter uncertainty and compute the variances using the predictive densities $p(R_{t+k}|y_t)$ in equation (5).

In practice the term structure of risk is computed numerically. We draw parameters $\theta^{(\ell)}$ from the posterior density $p(\theta|\mathbf{Y})$. For each draw $\theta^{(\ell)}$ ($\ell = 1, \dots, L$) we compute the conditional means and variances

$$E_{i\ell}(k) = E[R_{i,t+k}|\theta^{(\ell)}, y_t] \quad (7)$$

$$V_{i\ell}(k) = \text{var}[R_{i,t+k}|\theta^{(\ell)}, y_t] \quad (8)$$

using the formulas in appendix A, and where the conditional mean is evaluated at some fixed value y_t .¹ The term structure of risk is obtained by averaging over the parameters using the standard formula

$$\bar{V}_i(k) = \frac{1}{L} \sum_{\ell} V_{i\ell}(k) + \frac{1}{L} \sum_{\ell} (E_{i\ell}(k) - \bar{E}_i(k))^2 \quad (9)$$

where

$$\bar{E}_i(k) = \frac{1}{L} \sum_{\ell} E_{i\ell}(k)$$

The second term in equation (9) is the variance of the conditional mean and adds positively to the overall risk of each asset class. The Bayesian term structure of risk will therefore generally be larger than the conditional variance reported by Campbell and Viceira (2005). The term structure of risk also depends on the prior. More informative priors lead to less mean uncertainty and a smaller second term. The effect of different informative priors on the first term can go in both directions, depending on the specification of the prior.

Parameter uncertainty will have very limited effects for the short-term risk estimates. For large k , however, the uncertainty about the conditional mean will eventually become the dominating component. This can be seen most easily in the simple case without any dynamics ($\mathbf{B} = 0$). For a vector of risky assets we then have the returns

$$r_t = \mu + \varepsilon_t \quad (10)$$

with ε_t white noise with zero mean and known variance Σ . Parameter uncertainty in μ is summarized by the posterior

¹ In a linear vector autoregression the conditional mean $E_i[R_{t+k}]$ depends on y_t , which leads to state dependence in the term structure of risk. We find that the effect is very limited within our VAR model. We ignore the state dependence in the rest of the paper and evaluate all risk-related quantities at the unconditional mean.

$$p(\mu|\mathbf{Y}) \sim N(\hat{\mu}, \Omega) \quad (11)$$

Conditional on $\hat{\mu}$ the term structure is flat at

$$\frac{1}{k} \text{var}[R_{t+k}|\hat{\mu}] = \Sigma \quad (12)$$

Since cumulative returns

$$R_{t+k} = k\mu + \sum_{j=1}^k \varepsilon_{t+j} = k\hat{\mu} + k(\mu - \hat{\mu}) + \sum_{j=1}^k \varepsilon_{t+j} \quad (13)$$

also contain the cumulative parameter uncertainty $k(\mu - \hat{\mu})$, the unconditional term structure is

$$\frac{1}{k} \text{var}[R_{t+k}] = \Sigma + k\Omega \quad (14)$$

which is linearly increasing in k . The long-run risk increases, since estimation errors are perfectly correlated over time and thus accumulate quickly. As a benchmark case we evaluate the term structure at a horizon of $k=15$ years with expected returns estimated from a sample of $T=60$ years. With an uninformative prior we get $\Omega = \frac{1}{T}\Sigma$ and the second term is $\frac{k}{T}\Sigma = \frac{1}{4}\Sigma$ and raises the long-run risk by about 25%. Since uncertainty in the unconditional mean has such a profound effect on long-term risk and since this uncertainty could also be substantial, it is the focus of our econometric approach.

In this simple case without any dynamics the effect is proportionally the same for all asset classes, meaning that it will not be necessarily true that parameter uncertainty will put equity at a disadvantage relative to other asset classes. The three asset classes we consider have very different dynamics and therefore different term structures of risk. For T-bills the effect of parameter uncertainty is related to the persistence of inflation and interest rate shocks. For equity the shape of the term structure depends on the parameters that generate mean reversion. For bonds it is the interaction with the yield spread and inflation risk that are the main drivers. It will take the full VAR system to evaluate the effects of parameter uncertainty on all variances and covariances.

For assets with mean reversion the term structure will be decreasing for small k , exactly as in Campbell and Viceira (2005), since the effect of parameter uncertainty is usually very small for short horizons. Once there is parameter uncertainty, however, every term structure of risk will eventually be upward sloping beyond a certain horizon.

3. ASSET ALLOCATION

We derive the optimal portfolio for a power-utility investor who maximizes end-of-period wealth. The investor plans to hold constant proportions of his wealth in each of the asset classes until the end of the period. Fixed or stable portfolio weights appear closely connected to the industry practice of large institutional investors. Pension funds commonly plan their strategic portfolio on a constant mix basis and then allow various tactical bets depending on short-term market views. At time t the investor allocates wealth to real T-bills, stocks and bonds with portfolio weights $w = (w_r, w_s, w_b)$ and intends to keep the weights fixed until the end of the planning horizon at time $t+k$. Assuming power utility with risk aversion γ , the investor solves

$$\max_w E_t[U(W_{t+k})] = \max_w E_t \left[\frac{W_{t+k}^{1-\gamma}}{1-\gamma} \right] \quad (15)$$

When investment allocations are rebalanced back to the initial weights at the end of each period, the weights w_i remain constant and final wealth is given by

$$W_{t+k} = \prod_{j=1}^k \left(\sum_{i=r,s,b} w_i e^{r_{i,t+j}} \right) \quad (16)$$

where initial wealth is normalized at $W_t = 1$.

In order to calculate the maximum expected utility in equation (15) we need the distribution of future asset returns. As for the term structure of risk we first draw a set of parameters from the posterior distribution. Conditional on parameters $\theta^{(\ell)}$ we simulate a scenario of future returns $r_{t+j}^{(\ell)}$ from the VAR. Letting $W_{t+k}^{(\ell)}(w)$ denote the final wealth in scenario ℓ using portfolio weights w , we estimate the expected utility of a portfolio as the average realised utility over all scenarios:

$$\bar{U}_k(w) = \frac{1}{L} \sum_{\ell} U \left(W_{t+k}^{(\ell)}(w) \right) \quad (17)$$

The optimal portfolio is the w that maximizes expected utility in equation (17). We assume that short-sell constraints restrict the weights to be non-negative for all three asset classes, and that the weights sum to one.

The expectation in equation (15) may not always exist. Although we assume that the innovations in the VAR are normally distributed conditional on the parameters, the predictive densities (5) will have fatter tails than the normal distribution. If $\gamma > 1$, and if parameter uncertainty leads to fat tails and tail correlation for all assets, there may be too much probability mass on wealth being very close to zero and hence on utility values diverging to minus infinity. In our setting we cannot rule this out, since utility is defined on real returns that share inflation as a common factor. Infinite inflation is the event that leads to zero real wealth. The problem with expected utility with parameter uncertainty is well known (see, for example, Kandel and Stambaugh, 1996), and it is usually assumed that a risk-free asset exists. With positive investment in a risk-free asset there always exists a portfolio with finite utility. We do not have a real risk-free asset, but we can guarantee that portfolios have finite utility if we assume that the real return on nominal T-bills is always above a lower limit r_{\min} . In simulating paths of asset returns we reject draws that would violate the constraint $r_{\text{tb},t} > r_{\min}$.² If the T-bill has a positive weight in the portfolio, wealth will never go to zero and expected utility is bounded away from minus infinity. This approach ensures that the maximization problem (15) is well defined.

To gauge the order of magnitude of effects of parameter uncertainty on asset allocation, consider an investor who can choose from a menu of risky assets with logarithmic return vector r_t with mean μ and covariance matrix Σ . To avoid the numerical problems noted above, also consider the log-linear approximation of Campbell and Viceira (2002) for the portfolio returns:

$$r_{p,t} = w' \left(r_t + \frac{1}{2} \sigma^2 \right) - \frac{1}{2} w' \Sigma w \quad (18)$$

² In practice, we set the lower bound to 12 standard deviations below the sample mean, meaning that the probability of ever hitting this lower bound is virtually zero.

where $\sigma^2 = \text{diag}(\Sigma)$. Under the assumptions from Section 2, namely normality as in equation (10) and parameter uncertainty as in equation (11), the log of expected utility can be written as the mean–variance objective:

$$\frac{\ln((1 - \gamma)E[U(W_{t+k})])}{1 - \gamma} = k \left(w' \left(\hat{\mu} + \frac{1}{2}\sigma^2 \right) - \frac{1}{2}\gamma w' \left(\Sigma + k \left(1 - \frac{1}{\gamma} \right) \Omega \right) w \right) \tag{19}$$

The log-linear log-normal approximation leads to a standard mean–variance problem. With estimation error the covariance matrix Σ has an additional term $k(1 - 1/\gamma)\Omega$, which for $\gamma \neq 1$ depends on the investment horizon k and is positive definite for $\gamma > 1$. For $\Omega = \Sigma/T$, the optimal portfolio weights are

$$w_k = w^* + \left(1 - \iota' w^* \right) w_{\min} \tag{20}$$

with $w_{\min} = \Sigma^{-1} \iota / \iota' \Sigma^{-1} \iota$ the minimum variance portfolio and w^* the unrestricted mean–variance portfolio:

$$w^* = \frac{1}{1 + \frac{k}{T} \left(1 - \frac{1}{\gamma} \right)} \frac{1}{\gamma} \Sigma^{-1} \left(\hat{\mu} + \frac{1}{2}\sigma^2 \right) \tag{21}$$

For $\gamma = 5$ and $k/T = 15/60$ the adjustment factor in w^* is equivalent to an increase of the covariance matrix in equation (19) by 20% or an increase in risk aversion from $\gamma = 5$ to $\gamma^* = 6$. Since portfolio weights must sum to one, this will reduce the weight of the more risky assets in the portfolio. Estimation error shifts the portfolio allocation towards less risky assets. The portfolio depends on the horizon k ; the larger k , the more the portfolio moves towards the minimum variance portfolio.

In the full VAR model the returns on equity, long-term bonds and T-bills have rich dynamics and distinct term structures of risk. In the empirical analysis we do not work with the approximate utility (19) but with the simulated expected utility (17).

4. ECONOMETRIC SPECIFICATION

Parameter uncertainty is accounted for by a Bayesian analysis of the VAR system (3). We consider a series of alternative priors on the unconditional mean of the asset returns and state variables. For our econometric analysis we formulate informative priors on μ as

$$\mu \sim N \left(\mu_0, \frac{1}{\kappa} \Omega_0 \right) \tag{22}$$

where both μ_0 and Ω_0 are exogenously specified. The scalar parameter κ is a shrinkage factor. It represents the investor’s degree of confidence in the prior information. A shrinkage factor close to zero corresponds to a dispersed prior on μ . A large shrinkage factor gives much weight to the prior information. When $\kappa \rightarrow 0$ we obtain a non-informative prior.

We specify a flat prior on \mathbf{B} in the stationary region and an uninformative inverted Wishart on Σ :

$$p(\mathbf{B}, \Sigma) \propto \mathbb{I}(\mathbf{B}) \times |\Sigma|^{-(n+1)/2} \tag{23}$$

where the indicator function $\mathbb{I}(\mathbf{B})$ is equal to one if the maximum eigenvalue of \mathbf{B} is less than one, and zero otherwise. Like Barberis (2000) and Hollifield *et al.* (2003) we exclude non-stationary models,

since unconditional means cannot be defined in such models and because explosive roots also lead to explosive behaviour of the term structure of risk.³ When the likelihood attaches negligible probability mass to the non-stationary region, the posterior mode coincides with the least squares estimate of the parameters μ and \mathbf{B} .

The different priors for μ affect the estimates of the term structure of risk in two ways. First, even though μ and \mathbf{B} have independent priors, the posterior for \mathbf{B} is affected by the choice of prior for μ . This way, the prior for μ affects the long-run risk estimates through the first term of equation (9) in Section 2. Second, uncertainty in the expected returns is a separate source of risk, leading to the second term in equation (9).

For various reasons we explicitly parametrize the VAR using the unconditional mean μ instead of a linear constant $c = (I - \mathbf{B})\mu$. First, from an economic perspective it is more natural to formulate priors directly on the long-term outlook μ instead of the awkward nonlinear function c . This is especially relevant in our application, as informative priors on the level of future returns are already in place for some time in asset allocation at many institutional investors and asset managers (see Black and Litterman, 1992). This practice has been motivated by the extreme sensitivity of optimized portfolios with respect to small changes in expected returns. It is therefore common practice for long-term investors to base their future return expectations not only on historical data but also on current market circumstances, economic theory, human judgement and experiences in other countries.⁴ Historical data are far more accurate in providing estimates of future volatility and correlations. For our model this means a flat prior on \mathbf{B} and Σ , and an informative prior on the long-term economic outlook.

Second, the parametrization mitigates the econometric inconvenience of censoring the prior on the stationary region $\mathbb{I}(\mathbf{B})$ in the benchmark flat prior on c and \mathbf{B} . Villani (2005, 2009) shows that even a weakly informative prior on the unconditional mean of very persistent state variables reduces the probability of drawing explosive roots for the VAR. The parametrization with μ and \mathbf{B} in equation (3) allows efficient numerical computation of the posterior density. Villani 2009 shows that a simple three-block Gibbs sampler, described in appendix B, works very well with persistent time series.

To investigate the sensitivity of our long-run risk estimates we implement both optimistic outlooks and more negative views on the future. We want our alternative priors to be sufficiently distinct so as to generate meaningful model uncertainty. Specifically, we consider two different prior means μ_0 . For our application to a postwar US sample we split the historical data into two parts: NBER expansion periods and NBER contraction periods. Averages in the expansion periods represent a positive outlook for all asset returns and state variables in y_t , whereas the contraction period averages define a pessimistic outlook for long-term means. The alternative vectors of prior means will be labelled *optimistic* and *pessimistic*.

We distinguish between very confident views and highly dispersed priors on the long-run expected returns. For the prior precision κ in equation (22) we consider three different values. To calibrate κ we take an empirical perspective as well and estimate Ω_0 as the long-run variance of our historical data:

$$\Omega_0 = \frac{1}{T} \sum_{\ell=-L}^L v_\ell \left(\frac{1}{T} \sum_t (y_t - \bar{y})(y_{t+\ell} - \bar{y})' \right) \quad (24)$$

³ Even very small probabilities of explosive roots will dominate the results. If there is substantial evidence against stationarity of the VAR we would rather include VECM specifications or structural breaks. Adding these features will only add to the long-run risk and further strengthen our conclusion that parameter uncertainty is an important component of long-run risk.

⁴ Some examples for each of these sources of prior information include forward rates as an example of the current market view about interest rates. An example of theory-motivated priors for the equity premium is the macro-finance literature on the equity premium starting with Mehra and Prescott (1985). Welch (2001) is an example of survey evidence of expected returns. Dimson *et al.* (2002) compare average bond and equity premiums of 16 countries over more than a century, which would also constitute an excellent data-based shrinkage prior for the means in any single country.

Table I. Summary statistics

	Full sample, quarterly				Annualized, simple		
	ave	SE	SD	AR(1)	Full	Pessimist	Optimist
i	1.23	(0.28)	1.35	0.93	4.92	5.23	4.86
r_{tb}	0.32	(0.15)	1.33	0.46	1.26	1.47	1.22
d	-3.49	(0.20)	0.41	0.98	-3.49	-3.27	-3.53
S	0.30	(0.05)	0.59	0.80	1.21	1.33	1.19
x_s	1.18	(0.58)	14.91	0.13	5.83	1.88	7.31
x_b	0.16	(0.24)	7.83	-0.02	0.96	6.66	-0.15

Note: The first four columns report means ('ave'), standard deviations ('SD') and autocorrelations ('AR(1)') for the entire sample (1952Q2–2008Q4). Summary statistics are based on logarithmic returns in percent per quarter. Standard errors of the mean ('SE') are computed using the Newey–West estimate of the long-run variance. The final three columns report the implied annualized means for simple returns by adding half the variance and multiplying by four. The means are reported for both the full sample as well as for the two subsamples based on NBER business cycles that are used as optimist and pessimist priors. Variables are real 3-month T-bill return (r_{tb}), excess stock returns (x_s), excess bond returns (x_b), nominal Treasury bill return (i), log dividend yield (d) and term spread (S).

with Newey–West weights $v_\ell = (1 - \frac{\ell}{L+1})$. With this choice of Ω_0 we can interpret κ as the weight of the prior. We vary κ to increase or decrease the precision of the prior while keeping Ω_0 fixed. The prior precision parameter varies between $\kappa = 0.01$ (*uninformative*), $\kappa = 1$ (*moderate*) and $\kappa = 100$ (*dogmatic*). For $\kappa = 100$ we almost impose the unconditional means and all uncertainty is in the dynamic aspects of the term structure of risk and the hedge demands in asset allocation. If the prior precision factor $\kappa = 0.01$, the prior hardly carries any weight and all evidence will be data based.

5. ESTIMATION RESULTS

We consider returns on three asset classes (stocks, bonds and T-bills) and three state variables that help predict asset returns (inflation, dividend yield and term spread). For our empirical analysis we use quarterly US data. All series start in 1952Q1 and end in 2008Q4. The 90-day T-bill and the 10-year constant maturity yield are from the FRED website.⁵ In order to generate the yield spread we obtain the zero yield data from Duffee (2002).⁶ As these data are only available until 1998Q4, we have extended the series using the data from Gürkaynak *et al.* (2007).⁷ For inflation we use the non-seasonally adjusted consumer price index for all urban consumers and all items also from the FRED website. Data on stock returns and the dividend price ratio are based on the S&P Composite and are from the 'Irrational Exuberance' data of Shiller.⁸ We construct the gross bond return series from 10-year constant maturity yields using the log-linear approximation approach in Campbell *et al.* (1997).

Table I provides summary statistics. Even with the financial crisis year 2008 included, the sample equity premium of around 6% is larger than most recent studies on the prospective equity premium suggest. For example, Claus and Thomas (2001) suggest a forward-looking equity premium of about 3.5%—the same number that emerges from the average of 10 years of survey evidence in Graham and Harvey (2010). Since our post World War II sample period is rather short, the sample means of the equity and bond premia may be very poor estimates of the long-run expected returns. The standard errors of the annualized sample means are 1% and 2% for the bond and equity premium, respectively.

⁵ <http://research.stlouisfed.org/fred2/>.

⁶ <http://faculty.haas.berkeley.edu/duffee/affine.htm>.

⁷ <http://www.federalreserve.gov/econresdata/researchdata.htm>.

⁸ <http://aida.econ.yale.edu/shiller/data.htm>.

The final columns in Table I give the summary statistics for the prior means. We assign each observation in our sample period to NBER contraction or expansion periods. A contraction starts at the peak of a business cycle and ends at the trough, and the expansion vice versa. Nine contraction periods exist in our data sample, which have a duration of between two and six quarters. Ten expansion periods exist in our sample period, with a duration ranging from 4 to 40 quarters. Contractions are on average much shorter than expansions, and consequently 191 out of the 228 observations are assigned to expansions, and the remaining 37 observations are contractions.

The prior mean is set at either the average over the contraction or expansion subsample. The subsample averages support a wide range of prior expectations. Bonds seem less attractive than stocks during expansions, whereas they seem more attractive during contractions. Even though negative expected returns may be unreasonable from an asset pricing perspective, we leave the averages as they are, since they generate much dispersion in the priors. For our results on the term structure of risk we will look at each of the priors in isolation for their long-run risk properties. We can thus study the robustness of the dynamic properties with respect to the choice of the prior. For our portfolio results the dogmatic prior provides the opportunity to check the effect of extreme, perhaps implausible, priors on the portfolio choice.

The averages of state variables as the short interest rate, dividend yield, and term spread also differ between the two periods. Fama and French (1989) link the dividend yield and yield spread to the business cycle. They argue that the risk premia are high in contraction periods and low in expansion periods. The opposite applies to the dividend price ratio which is high in expansion periods and low in contraction periods.

The VAR system has been estimated on the entire sample. First, we verified that a first-order VAR is a reasonable choice. Based on the Schwartz criterion it is the preferred model among lag lengths ranging from zero to four. Tables II and III summarize the ordinary least squares (OLS) parameter estimates

Table II. Parameter estimates

	i_t	$r_{tb,t}$	d_t	S_t	$x_{s,t}$	$x_{b,t}$	R^2
i_{t+1}	0.96 (0.03)	0.00 (0.03)	0.04 (0.04)	0.08 (0.06)	0.00 (0.00)	0.00 (0.00)	0.89
$r_{tb,t+1}$	0.21 (0.07)	0.46 (0.07)	-0.03 (0.10)	0.27 (0.15)	0.00 (0.01)	-0.01 (0.01)	0.27
d_{t+1}	0.01 (0.01)	-0.01 (0.01)	0.97 (0.01)	0.00 (0.02)	0.00 (0.00)	0.00 (0.00)	0.97
S_{t+1}	0.02 (0.02)	-0.01 (0.02)	-0.03 (0.03)	0.82 (0.05)	0.00 (0.00)	0.00 (0.00)	0.66
$x_{s,t+1}$	-1.58 (0.91)	0.95 (0.88)	3.64 (1.28)	0.10 (1.88)	0.11 (0.07)	0.25 (0.14)	0.09
$x_{b,t+1}$	0.46 (0.48)	0.39 (0.47)	-0.42 (0.68)	3.32 (0.99)	-0.08 (0.04)	-0.09 (0.07)	0.07

Note: The table reports OLS parameter estimates of the VAR $y_{t+1} = \mu + \mathbf{B}(y_t - \mu) + \varepsilon_{t+1}$ with variables: real 3-month T-bill return (r_{tb}), excess stock returns (x_s), excess bond returns (x_b), nominal Treasury bill return (i), dividend yield (d) and term spread (S). Standard errors are in parentheses. The last column contains the R^2 .

Table III. Residual correlation matrix

	i	r_{tb}	d	S	x_s	x_b
Nominal T-bill rate (i)	0.23	—	—	—	—	—
Real T-bill rate (r_{tb})	-0.34	0.58	—	—	—	—
Log dividend yield (d)	0.05	-0.08	0.07	—	—	—
Term spread (S)	-0.83	0.16	-0.03	0.18	—	—
Excess stock returns ($x_{s,t}$)	-0.03	0.06	-0.98	0.02	7.23	—
Excess bond returns ($x_{b,t}$)	-0.65	0.40	0.01	0.14	-0.02	3.83

Note: The table reports the residual correlation matrix $\hat{\Sigma}$ for the OLS estimates of the VAR $y_{t+1} = \mu + \mathbf{B}(y_t - \mu) + \varepsilon_{t+1}$. Diagonal entries are standard deviations; off-diagonal entries are correlations.

together with the correlations and standard deviations of the residuals. Since our model is similar to much of the literature, except for the sample period, we only highlight the most important results. First, the three state variables (nominal interest rate, dividend yield, term spread) are almost univariate AR(1) processes. Second, as in Ang and Bekaert (2007), the nominal interest rate and the dividend price ratio jointly predict excess stocks returns. As in Campbell and Viceira (2005), the combination of a negative correlation of shocks to the dividend price ratio and stocks, and the positive predictive coefficient of the dividend price ratio, imply mean reversion in stocks returns. The excess return on bonds is related to the yield spread, the nominal interest rate and stock returns. Third, bond returns are also mean-reverting. The nominal interest rate is a predictor of excess bond returns, which has the required opposite signs of the predictive coefficient and residual correlation. The term spread leads to a mean aversion part. The R^2 around 8% for both stocks and bond returns is of the magnitude typically found for quarterly data. It is low enough to be credible for efficient financial markets and large enough to be economically meaningful at longer horizons (Campbell and Thompson, 2008).

The priors influence the persistence of state variables and the predictability of stock and bond returns. Table IV indicates that the posterior mean of the autocorrelation parameter of the two most persistent state variables—dividend yield and nominal T-bill rate—increases when we specify an uninformative prior on the long-run means μ . The differences are small, but potentially meaningful so close to the unit root. The largest differences are for the predictability parameter of the dividend price ratio and excess stocks returns: predictability varies from large and very significant under the uninformative prior to small and insignificant for the *dogmatic pessimist* prior. Predictability of bonds is less affected by the prior.

An important parameter for the long-term risk estimates is the maximum eigenvalue of the VAR system. From the OLS estimates of the coefficient matrix we have the point estimate $\lambda_{max} = 0.976$, implying that the VAR is stationary with a half-life of a little more than 7 years. Figure 1 shows the posterior distribution of λ_{max} for the different priors. Even though the priors are only formulated on the unconditional means, they have a strong impact on the persistence of the system. For all priors the posterior mean of λ_{max} is larger than the OLS estimate. The posterior for an *uninformative* prior on the unconditional mean is extremely skewed: the mode of the distribution is 0.996, almost at the unit root.⁹ The prior therefore mitigates the problem of a downward bias of autoregressive parameters in a VAR with near unit roots. The larger persistence leads to steep initial mean reversion in the term structure of risk, but also to a quick reversal towards an upward sloping term structure.

6. LONG-TERM RISKS

Figure 2 shows different estimates of the term structure of risk. Conditional on the OLS estimates of the VAR the results are very similar to the figures in Campbell and Viceira (2005) and others. Bills become more risky with the investment horizon due to roll-over risk and inflation risk. Stocks exhibit strong mean reversion that reduces the long-term risk by a factor two. Real returns on nominal bonds show some mean reversion as well, but far less than equity.

Adding parameter uncertainty increases the risk for all three asset classes. With the *uninformative* priors for the long-term mean, the term structure is hardly affected by parameter uncertainty at very short horizons. At longer horizons parameter uncertainty becomes an important risk component. For equity the mean reversion becomes much more limited: the annualized standard deviation moves from 15% at the one quarter horizon to 13.5% at the 15-year horizon.¹⁰ Most striking is the almost uniform effect on all three asset classes. At the 15-year horizon an *uninformative* prior implies a 25% higher

⁹ In the Monte Carlo simulations we reject about 10% of draws because of $\lambda_{max} > 1$. This truncates the right tail of the posterior of the maximum eigenvalue, but nevertheless leads to a posterior mean and mode that are above the OLS point estimate.

¹⁰ If we extend the horizon we find that the term structure of equity risk starts to increase from about the 50-year horizon.

Table IV. Posterior means of selected VAR parameters

Parameter	OLS	P_{01}	P_1	P_{100}	O_{01}	O_1	O_{100}
i_{t+1}, i_t	0.961 (0.028)	0.967 (0.026)	0.957 (0.026)	0.957 (0.026)	0.967 (0.026)	0.960 (0.027)	0.958 (0.027)
d_{t+1}, d_t	0.968 (0.013)	0.972 (0.013)	0.986 (0.009)	0.979 (0.010)	0.972 (0.013)	0.969 (0.012)	0.968 (0.013)
$x_{s,t+1}, i_t$	-1.58 (0.91)	-1.49 (0.88)	-1.11 (0.94)	-1.39 (0.97)	-1.50 (0.88)	-1.51 (0.90)	-1.55 (0.91)
$x_{s,t+1}, d_t$	3.64 (1.28)	3.30 (1.28)	1.67 (0.90)	0.82 (1.15)	3.28 (1.25)	3.60 (1.21)	3.53 (1.24)
$x_{b,t+1}, i_t$	0.46 (0.48)	0.35 (0.45)	0.35 (0.47)	0.47 (0.48)	0.37 (0.45)	0.48 (0.46)	0.52 (0.46)
$x_{b,t+1}, S_t$	3.32 (0.99)	3.17 (1.00)	2.97 (0.99)	3.73 (1.02)	3.19 (1.00)	3.29 (0.99)	3.39 (1.00)

Note: This table shows the effect of different priors on selected VAR parameters \mathbf{B}_{ij} . 'OLS' are the least squares estimates (standard errors). All other entries denote posterior means (posterior standard deviations). O = textitoptimist, P = *pessimist*, 100 = *dogmatic*, 01 = *uninformative*.

annualized standard deviation for bills, bonds as well as equity. Not only does equity become more risky, but all asset classes and all risk estimates increase proportionally.

Informative priors can lead to very different term structures of risk. The lines for *dogmatic optimist* and *dogmatic pessimist* show the effect of fixing the unconditional means, thereby leaving out one component of the risk estimates. For the *dogmatic optimist* prior this indeed leads to a reduction of the risk estimates compared to the *uninformative* prior, albeit that the effect is very small. For the *dogmatic pessimist* prior the effect is much stronger. Here the risk increases, because the prior not only affects inference on the long-run mean, but also the inference on the dynamics. In Section 5 we discussed the effect on the coefficient relating the predicted equity return to the dividend yield. This coefficient is smaller in the *dogmatic pessimist* prior and hence the mean reversion disappears. The *dogmatic* priors have less effect on the risk estimates for the other asset classes, since for these asset classes the prior has very little impact on the most important parameters for their return dynamics.

Figure 3 offers a closer look at the term structure of equity risk. The figure shows a decomposition of the total equity risk in the average conditional variance and the variance of the conditional mean, as in equation (9). The first term, the average conditional variance, is the biggest component. The big risk increase from parameter uncertainty is related to the concavity of the relation between the term structure of risk and the persistence of the system.¹¹ Evaluating the term structure at the point estimates is very different from the average term structure across many different draws from the posterior density of the parameters. The variance of the expected returns is negligible at short horizons; even at 50 years it is still much smaller than the first term.

Correlations are robust against parameter uncertainty. Figure 4 shows the pairwise correlations among the real returns. The correlations exhibit interesting patterns as a function of the investment horizon, but they are hardly affected by parameter uncertainty. For the *uninformative* priors the correlations are hardly distinguishable from the OLS results. Again, the impact of parameter uncertainty on the covariance is proportional to the impact on the variance, and cancels out in the correlations. The only exception is the *dogmatic pessimist* prior, which leads to a very different correlation between equity and bonds at all horizons. In this case the prior mean is far from the sample mean and as we noted before this has an impact on the dynamic properties of the system. The *dogmatic pessimist* prior leads to slightly negative correlation between stocks and bonds, and hence implies much more diversification benefits than the other priors.

As a descriptive measure for the inflation hedge properties of the three asset classes we examine the correlation of nominal returns and inflation over different horizons. Results in Figure 5 show that the inflation hedge qualities of T-bills, stocks and a constant maturity treasury portfolio are robust to

¹¹ See the online Appendix (supporting information) for a more detailed analysis.

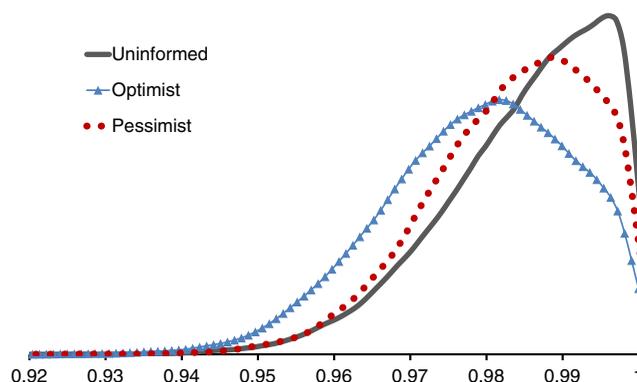


Figure 1. Posterior density of maximum eigenvalue. The figure shows the posterior density of the maximum eigenvalue λ_{max} in the VAR system. *Uninformed* refers to the priors with $\kappa=0.01$ (optimist and pessimist are indistinguishable). *Optimist* and *Pessimist* refer to the dogmatic priors with $\kappa=100$

incorporating parameter uncertainty. The effect of parameter uncertainty on the covariance between the nominal returns and inflation is proportional to the effect on the corresponding variances. Even the *dogmatic pessimist* prior does not lead to really different estimates. The horizon effects themselves are in line with Hoevenaars *et al.* (2008) for fixed parameter estimates. Rolling over 3-month T-bills ensures that the lagged inflation is incorporated, and consequently the T-bill is the best inflation hedge among the asset classes we consider at all investment horizons. At long horizons constant maturity bonds become an inflation hedge as well. However, due to the inverse relationship between yield changes and bond prices, the short-term inflation hedging properties are poor. For stocks there is hardly any correlation with inflation at all on horizons up to 15 years. For a positive correlation we would have to extend the horizon to more than 25 years. For the *dogmatic pessimist* prior the correlation remains negative at all horizons.

7. PORTFOLIO IMPLICATIONS

For the effects on optimal portfolio choice we first compare the conditional portfolios based on the OLS estimates with the Bayesian portfolios implied by the *uninformative* priors. Obviously, the alternative *dogmatic* priors shift the optimal portfolios in a predetermined direction as they condition on widely different long-term expected returns. (Table V)

Table VI contains the results for a benchmark case with risk aversion $\gamma=5$ and various horizons ranging from 1 to 15 years. For the OLS estimates we obtain the standard result that the weight of stocks increases with the investment horizon. This result is a combination of the high historical equity premium and the strong mean reversion conditional on the estimated VAR models. For horizons longer than 15 years the investor would prefer to be fully invested in stocks. T-bills have a high weight for short investment horizons, but reinvestment risk makes them less attractive at longer horizons. Due to the low bond premium in the data, bonds are not in the optimal portfolio.

As anticipated, parameter uncertainty does not matter much at the short 1-year horizon. For long horizons, however, the optimal investment in equity falls substantially. With the *uninformative* prior the optimal allocation to equity is always in the 50–60% range. Consequently, the horizon effect diminishes: the difference in allocation to equity between the 1-year and 15-year horizon is never more than 15%; whereas, as the share of equity decreases, the allocation to T-bills increases. T-bills are a good risk diversifier due to the very low correlation with stocks at longer horizons. Bonds are not in

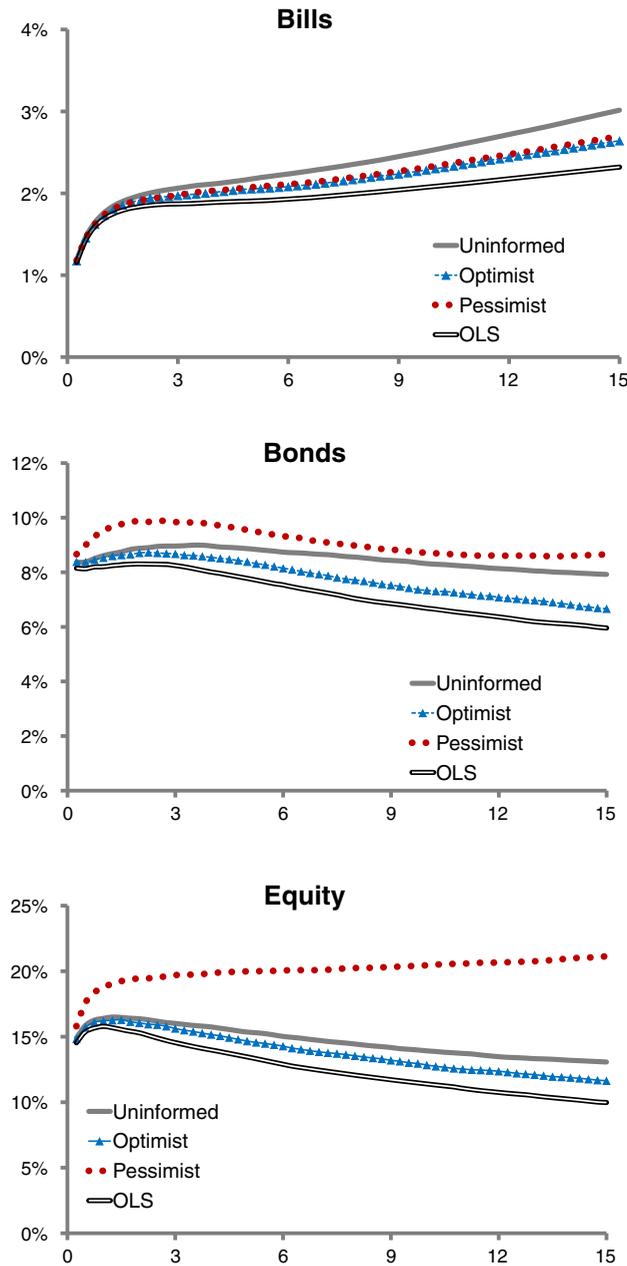


Figure 2. Term structure of risk and informative priors. The figure shows the term structures of risk for bills, bonds and equity. The estimates are either conditional on the OLS parameter estimates or based on the *uninformative* ($\kappa = 0.01$) or *dogmatic* ($\kappa = 100$) prior on the unconditional mean. For the *dogmatic* priors the figure shows both the *optimist* and the *pessimist* risk estimates. For the *uninformative* priors the two are indistinguishable. The horizontal axis denotes the horizon in years. Vertical axis is the annualized standard deviation of the predictive distribution

the optimal portfolio, because their risk rises in the same proportion as the risk of stocks, while the bond premium is much lower. Although ignoring parameter uncertainty leads to an overallocation to stocks at all horizons, the investor should still increase the weight of equity for investment horizons

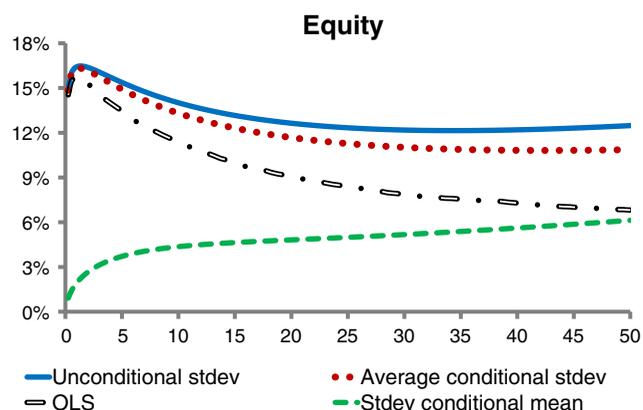


Figure 3. Decomposition of equity risk. The figure shows the term structure of risk for real equity returns. *OLS* denotes the term structure conditional on the OLS estimates. The other lines are derived from the predictive scenarios based on the *uninformative* prior. The horizontal axis denotes the horizon in years

up to 15 years. Parameter uncertainty causes mean reversion in equity largely to disappear, but since it also increases the risk of the other asset classes equity remains attractive for long-term investors.

Obviously, the optimal asset allocation varies substantially with the view and confidence the investor has about expected future asset returns. Under the *dogmatic pessimist* view long-term bonds have a high premium and therefore dominate the portfolio. Under the *dogmatic optimist* prior the short-term allocation to equity is much more than conditional on the OLS estimates. This is the optimistic mean effect. At long horizons the allocation is much lower than in the OLS case, since parameter uncertainty about the dynamics increases the risk.

Figure 6 shows how the risk attitude of the investor influences the long-term (15 years) optimal portfolio choice under the different priors. The pattern is the same in all cases, except for the *dogmatic pessimist* prior. Beyond a certain risk aversion, the allocation to stocks becomes less than 100% and negatively related to the risk aversion (γ). The trade-off is between equity and T-bills. Nominal bonds are only attractive under the *dogmatic pessimist* prior.

8. DISCUSSION

In our empirical analysis we have assumed a first-order VAR, used quarterly post World War II US data, and selected a flat prior on the dynamic properties of returns. Our model and data are motivated by the choices of Campbell and Viceira (2005) that we take as our point of departure for the formal Bayesian analysis of the effects of parameter uncertainty.

In this study we considered a limited number of alternative priors. We have not exhaustively looked at the many different dynamic models of asset returns and priors that exist in the literature, but rather focused on the mechanism by which parameter uncertainty increases risk at different horizons for different asset classes. Alternative priors may come up with lower risk, because they reduce parameter uncertainty by specifying an informative prior on the VAR dynamics. Wachter and Warusawitharana (2009), for example, specify a prior on the R^2 of a return prediction equation, while Pastor and Stambaugh (2009) add information on mean reversion through an informative prior on the correlation between shocks to expected and unexpected returns.

A first-order VAR is the most frequently used model for strategic asset allocation studies. Yet the model has limitations and risk estimates may increase by extending the class of models. For example, Guidolin and Timmermann (2007) and Pettenuzzo and Timmermann (2011) introduce structural

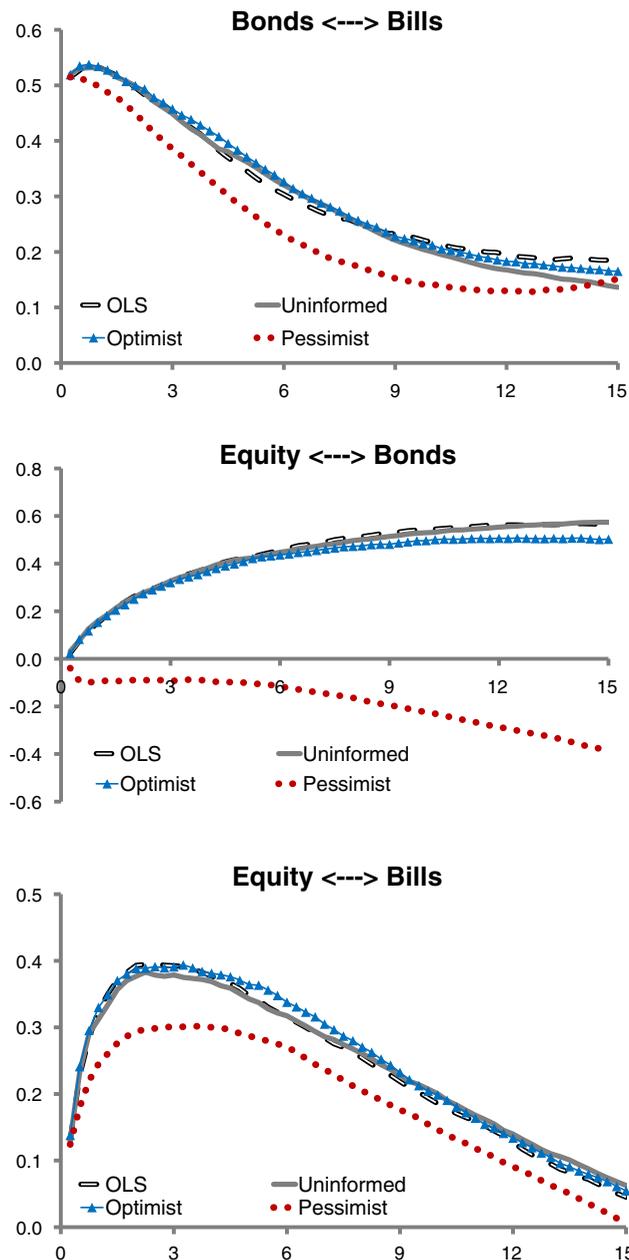


Figure 4. Correlations and prior information. The figure shows the effect of prior information on the correlation among real holding period returns of stocks, bonds and bills. *OLS* are the correlations conditional on the *OLS* estimates of the VAR. The other lines are Bayesian predictive correlations, based on the *uninformative* prior, *dogmatic* ($\kappa = 100$) *optimist* and *pessimist* priors on the unconditional means. The horizontal axis denotes the horizon in years

breaks and the risk of future structural breaks, while Pastor and Stambaugh (2012) warn against misspecification of the conditional mean. Avramov (2002) and Cremers (2002) suggest alternative additional prediction variables.

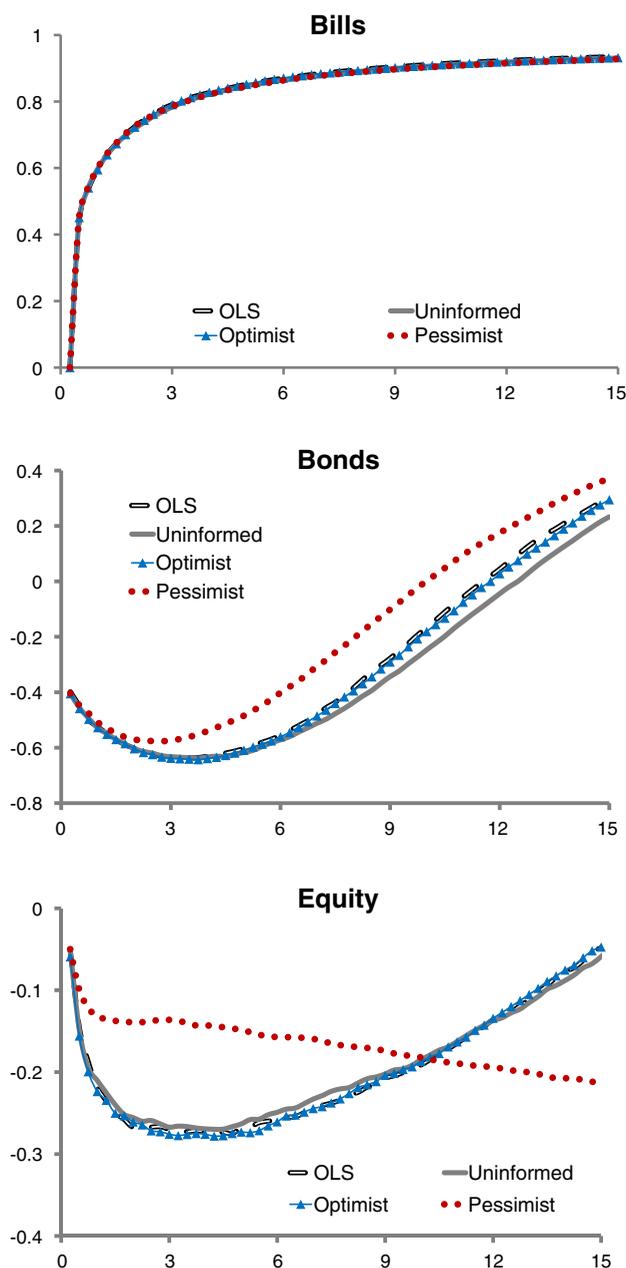


Figure 5. Inflation hedge. The figures show the correlation between shocks to cumulative nominal returns and shocks to cumulative inflation over different horizons. *OLS* is the correlation conditional on the parameter estimates. The Bayesian estimates are based on the *uninformative* prior ($\kappa = 100$) and the *optimist* and *pessimist* priors for the *dogmatic* design ($\kappa = 100$). On the horizontal axis is the horizon in years; on the vertical axis is the correlation

Even though a model with a single lag is the preferred model according to the Schwartz information criterion, it is worthwhile to check the sensitivity to lag length. For a second-order VAR we find the same term structure of risk conditional on the OLS estimates, and very similar portfolio allocations. Most important for the term structure of risk are the long-run dynamic properties of the VAR, which

Table V. Unconditional means

	Sample	VAR	P_{01}	P_{100}	O_{01}	O_{100}
Nominal rate	1.23 (0.28)	1.06 (0.45)	1.15 (0.73)	1.32	1.07 (0.73)	1.21
Real T-bill	0.32 (0.15)	0.27 (0.14)	0.29 (0.23)	0.37	0.28 (0.24)	0.29
Dividend yield	-3.49 (0.20)	-3.62 (0.23)	-3.52 (0.39)	-3.25	-3.63 (0.38)	-3.54
Term spread	0.30 (0.05)	0.32 (0.06)	0.30 (0.10)	0.33	0.32 (0.09)	0.30
Equity premium	1.18 (0.58)	0.93 (0.26)	0.93 (0.45)	-1.11	0.92 (0.45)	1.52
Bond premium	0.16 (0.24)	0.21 (0.09)	0.19 (0.14)	1.48	0.21 (0.15)	-0.08

Note: The table shows the effect of different priors on the unconditional means. 'Sample' denotes the sample average with the Newey–West standard error in parentheses; 'VAR' reports the nonlinear OLS estimates $\hat{\mu}$ with asymptotic standard errors in parentheses. The other entries denote posterior means (standard deviations) for the different priors (O = *optimist*, P = *pessimist*, $O1$ = *uninformative*, 100 = *dogmatic*). Posterior standard deviations for the dogmatic priors are zero. Units for returns are percent per quarter and refer to continuously compounded returns.

Table VI. Portfolio choice, investment horizon and priors

	OLS	P_{01}	P_1	P_{100}	O_{01}	O_1	O_{100}
<i>1-year horizon</i>							
Bills	0.28	0.35	0.20	0.01	0.35	0.45	0.37
Bonds	0.24	0.24	0.61	0.99	0.21	0.10	0.00
Equity	0.48	0.41	0.19	0.00	0.44	0.45	0.63
<i>5-year horizon</i>							
Bills	0.34	0.48	0.30	0.01	0.47	0.47	0.29
Bonds	0.03	0.09	0.55	0.99	0.07	0.00	0.00
Equity	0.63	0.43	0.15	0.00	0.46	0.53	0.71
<i>10-year horizon</i>							
Bills	0.20	0.48	0.24	0.01	0.47	0.39	0.22
Bonds	0.00	0.03	0.59	0.99	0.00	0.00	0.00
Equity	0.80	0.49	0.17	0.00	0.53	0.61	0.78
<i>15-year horizon</i>							
Bills	0.05	0.50	0.20	0.01	0.45	0.36	0.23
Bonds	0.00	0.00	0.61	0.99	0.00	0.00	0.00
Equity	0.95	0.50	0.19	0.00	0.55	0.64	0.77

Note: Optimal portfolio choice under power utility for an investor with a k -period investment horizon and quarterly rebalancing to a constant mix with risk aversion $\gamma = 5$ and for different investment horizons and different priors (O = *optimistic*, P = *pessimistic*, $O1$ = *uninformative*, 100 = *dogmatic*). Short-selling restrictions have been imposed. All results are based on simulations using 40,000 scenarios.

are not much affected by adding more lags. Bayesian results with the *uninformative* prior are shown in Figure 7. For short investment horizons, up to 2 years, the term structures are almost identical. For longer investment horizons the VAR(2) implies more risk for all asset classes. The larger risk estimates are due to the additional parameter uncertainty. A second lag adds 36 parameters to the system and with a flat prior this introduces substantial additional parameter uncertainty. As a consequence the upward-sloping part of the term structures starts at an earlier horizon. With more lags parameter uncertainty adds more and more risk unless we change the flat prior to an informative prior or a shrinkage prior that penalizes model complexity.¹²

A particular example of adding more dynamics to a VAR(1) is Pastor and Stambaugh (2012). They augment the prediction equation for equity returns by a latent component π_t that follows an independent univariate AR(1) process:

¹² Examples of such shrinkage priors are the Minnesota prior in Doan *et al.* (1984) or the reference priors for VARs in Ni and Sun (2003).

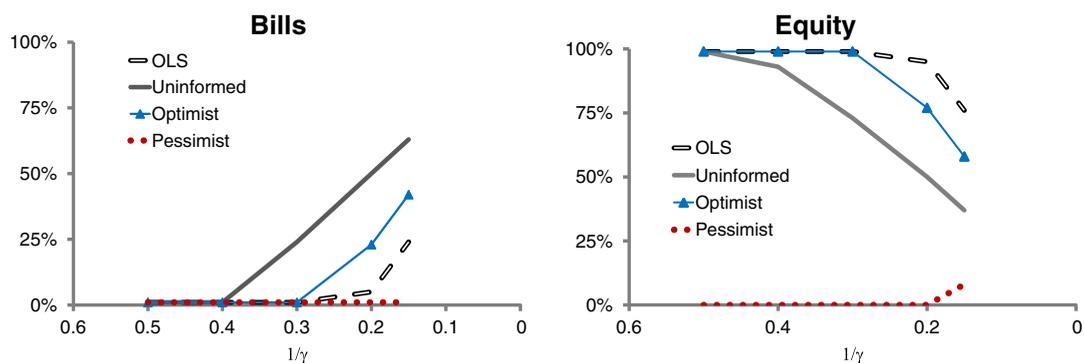


Figure 6. Portfolio holdings. The figure shows the optimal asset allocation for the OLS estimates and for different Bayesian priors as a function of the risk aversion parameter γ . The investment horizon is 15 years. Since weights sum to one, the allocation to bonds is one minus the weights for bills and equity shown in the figure

$$r_{s,t+1} = c + b'_z z_t + \pi_t + \varepsilon_{s,t+1} \tag{25}$$

$$\pi_{t+1} = \delta \pi_t + \eta_{t+1} \tag{26}$$

which is equivalent to

$$r_{s,t+1} = (1 - \delta)c + \delta r_{s,t} + b'_z z_t - \delta b'_z z_{t-1} + \varepsilon_{s,t+1} - \delta \varepsilon_{s,t} + \eta_t \tag{27}$$

Compared to a VAR(1) the prediction equation contains a restricted second lag of the predictor state variables z_{t-1} and an MA(1) error. It is therefore a specific parsimonious form of dynamic misspecification of the VAR(1). With uninformative priors the term structure of risk would not be very different compared to our Bayesian estimates of a VAR(1) or VAR(2). Pastor and Stambaugh (2012), however, specify an informative prior on δ , σ_η and the covariance between η_t and $\varepsilon_{s,t}$. When the prior mean of σ_η is centred away from zero, these informative priors imply an upward-sloping term structure for equity. With a prior close to zero, their results for quarterly data are comparable to our VAR.¹³

Pastor and Stambaugh (2012) attribute their high-risk estimates for equity to ‘imperfect predictors’. The term refers to omitted variables in the return prediction equations. Whereas they infer evidence of omitted predictors from lagged return data, there is a rich literature searching for prediction variables for equity returns. In our VAR we followed Campbell and Viceira (2005) and only included the dividend yield, nominal interest rate, inflation and the term spread as predictors in order to obtain a low-dimensional VAR system. For equity prediction Cremers (2002) and Avramov (2002) identify additional predictors like the credit spread and liquidity using a Bayesian model selection procedure. For bonds Cochrane and Piazzesi (2005) show that a linear combination of forward rates is a powerful predictor. For T-bills the monetary economics literature adds real output as a predictor. How inclusion of these variables affects the term structure of risk depends on their own dynamics and the covariance of their shocks with return shocks. This requires a much larger-dimensional VAR system. Diris (2011) develops new benchmark priors and algorithms to deal with risk estimates in such large-scale VAR models.

¹³ See the online Appendix (supporting information) for a closer comparison of results.

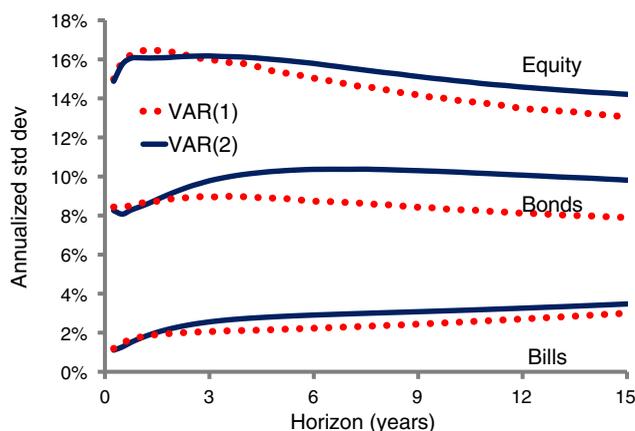


Figure 7. VAR(2). The figure compares the term structures of risks implied by a VAR(1) and VAR(2) using an uninformative prior ($\kappa=0.01$) on the long-run means and a flat prior on all dynamic parameters

The length of the sample period is also important for the effect of parameter uncertainty. Our estimates are based on about 60 years of post World War II data, whereas other studies work with data starting in 1926 or the much longer Shiller data starting in 1872. Samples that include the pre World War II data typically lead to more mean reversion for equity. A longer sample will generally result in less parameter estimation error if we stay within the same first-order VAR. But the longer the sample, the more urgent the question if the model structure and model parameters remain constant over time. The longer the sample period, the more likely there will be structural breaks or other forms of parameter instability.

Regime-switching models are a popular way to allow for structural breaks. Most importantly, these models open the possibility of future regime switches to incorporate the effects of model uncertainty and parameter changes on long-term risk estimates. An example is Guidolin and Timmermann (2007), who estimate the long-run volatility using a VAR(1) with separate parameters in each of four different regimes that follow a Markov switching process. They compare the term structures for large cap stocks and bonds between a regular VAR(1) and a regime-switching VAR(1). Conditional on the estimated parameters the regime-switching model implies very similar term structures of risk as a first-order VAR estimated on the full sample.¹⁴ The mean reversion component due to the interaction of the slow variation in the dividend yield with its negative covariance with equity returns is not too sensitive to the nonlinearities induced by the different regimes. The VAR(1) seems to capture the long-run properties of the data. Another effect in the regime-switching models is heteroskedasticity with different innovation variances in each regime. Starting in a low-volatility regime, the term structure will initially be upward sloping, whereas it will be downward sloping conditional on a high-volatility initial state. Unconditionally, i.e. without knowing in which regime we are, the heteroskedasticity effect has a very limited effect on the term structure of risk at longer horizons.

Guidolin and Timmermann (2007) do not consider parameter uncertainty, so we can only conjecture on the effects. Since the regime-switching model has separate parameters for each regime, the model contains more than four times as many parameters as a VAR(1). With an uninformative prior this will result in a large increase in the term structure of risk at longer horizons for all asset classes, similar to our experience with a second-order VAR. Again, one would need a form of informative priors to arrive at plausible estimates of the long-run volatility.

¹⁴ See Figure 6 in Guidolin and Timmermann (2007).

9. CONCLUSION

We studied the impact of parameter uncertainty on long-term risk and asset allocation of long-term investors who can invest in stocks, bonds and T-bills. When using uninformative priors, we find that real returns on stocks and bonds exhibit limited mean reversion at short and medium investment horizons. Mean reversion disappears at long horizons. Parameter uncertainty affects all three asset classes (stocks, bonds and T-bills) almost proportionally, with annualized standard deviations of returns increasing by a factor 1.25 at the 15-year horizon relative to estimates conditional on maximum likelihood parameter estimates. Asset return correlations are stable under weakly and moderately informative priors. Similarly, the correlations between inflation and nominal asset returns are also robust to incorporating parameter uncertainty.

We also considered informative priors on the long-run mean of returns and predictor variables and find that for highly persistent time series (like dividend yield and nominal interest rate) prior information changes the estimated persistence of shocks and the predictability of excess returns. Both effects increase the term structure of annualized volatility. A strong, almost dogmatic, pessimistic prior view on long-run stock and bond returns also affects the correlation between returns on stocks and bonds. Our results appear robust with respect to extensions of the class of models beyond a first-order VAR. Regime-switching models, for example, produce very similar unconditional term structures of risk.

Parameter uncertainty alters optimal portfolio allocations. Obviously, assets with high prior expected returns get a large portfolio weight, but more interesting is the effect of an uninformative prior. Even though the risk of all asset classes increases, the long-term portfolio weight of equity decreases substantially, and the weight of short-term bonds increases. Moreover, the effect of the investment horizon on the optimal portfolio diminishes.

ACKNOWLEDGEMENTS

We thank Michael Brandt, Ralph Koijen, Franz Palm and Luis Viceira for constructive comments on an earlier version, which was presented at the Econometric Society European Meeting, European Quantitative Forum in London, a Netspar pension workshop, and the CREATES workshop on Dynamic Asset Allocation in Aarhus. We thank participants for their comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of our employer or our colleagues.

REFERENCES

- Ang A, Bekaert G. 2007. Stock return predictability: is it there? *Review of Financial Studies* **20**: 651–707.
- Avramov D. 2002. Stock return predictability and model uncertainty. *Journal of Financial Economics* **64**: 423–458.
- Barberis N. 2000. Investing for the long run when returns are predictable. *Journal of Finance* **55**: 389–406.
- Bauwens L, Lubrano LM, Richard JF. 2003. *Bayesian Inference in Dynamic Econometric Models*. Oxford University Press: New York.
- Black F, Litterman R. 1992. Global portfolio optimization. *Financial Analysts Journal* **48**: 28–43.
- Brandt MW, Goyal A, Santa-Clara P, Stroud JR. 2005. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies* **18**: 831–873.
- Campbell JY, Thompson S. 2008. Predicting the equity premium out of sample: can anything beat the historical average? *Review of Financial Studies* **21**: 1509–1531.
- Campbell JY, Viceira LM. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press: Oxford.
- Campbell JY, Viceira LM. 2005. The term structure of the risk–return tradeoff. *Financial Analysts Journal* **61**: 34–44.
- Campbell JY, Chan YL, Viceira LM. 2003. A multivariate model for strategic asset allocation. *Journal of Financial Economics* **67**: 41–80.

- Campbell JY, Lo AW, MacKinlay AC. 1997. *The Econometrics of Financial Markets*. Princeton University Press: Princeton, NJ.
- Claus J, Thomas J. 2001. Equity premia as low as three percent? Evidence from analysts' earnings forecasts for domestic and international stocks markets. *Journal of Finance* **56**: 1629–1666.
- Cremers MKJ. 2002. Stock return predictability: a Bayesian model selection perspective. *Review of Financial Studies* **15**: 1223–1249.
- Cochrane JH, Piazzesi M. 2005. Bond risk premia. *American Economic Review* **95**: 138–160.
- Dangl T, Halling M. 2012. Predictive regressions with time-varying coefficients. *Journal of Financial Economics* **106**: 157–181.
- Dimson E, Marsh P, Staunton M. 2002. *Triumph of the Optimists*. Princeton University Press: Princeton, NJ.
- Diris B. 2011. Model uncertainty for long-term investors. SSRN Working Paper 1786587.
- Doan T, Litterman RB, Sims CA. 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews* **3**: 1–100.
- Duffee GR. 2002. Term premia and interest rate forecasts in affine models. *Journal of Finance* **57**: 405–443.
- Fama EF, French KR. 1989. Business conditions and the expected returns on stocks and bonds. *Journal of Financial Economics* **25**: 23–49.
- Graham JR, Harvey CR. 2010. The equity risk premium in 2010. SSRN Working Paper 1654026.
- Guidolin M, Timmermann A. 2007. Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control* **31**: 3503–3544.
- Gürkaynak RS, Sack B, Wright JH. 2007. The U.S. treasury yield curve: 1961 to the present. *Journal of Monetary Economics* **54**: 2291–2304.
- Hoevenaars RPMM, Molenaar RDJ, Schotman PC, Steenkamp TBM. 2008. Strategic asset allocation with liabilities: beyond stocks and bonds. *Journal of Economic Dynamics and Control* **32**: 2939–2970.
- Hollifield B, Koop G, Li K. 2003. A Bayesian analysis of a variance decomposition for stock returns. *Journal of Empirical Finance* **10**: 583–601.
- Johannes M, Korteweg A, Polson N. 2012. Sequential learning, predictive regressions, and optimal portfolio returns. SSRN Working Paper 1108905.
- Kandel S, Stambaugh R. 1996. On the predictability of stock returns: an asset-allocation perspective. *Journal of Finance* **51**: 385–424.
- Mehra R, Prescott EC. 1985. The equity premium puzzle. *Journal of Monetary Economics* **15**: 145–161.
- Ni S, Sun D. 2003. Noninformative priors and frequentist risks of Bayesian estimators of vector-autoregressive models. *Journal of Econometrics* **115**: 159–197.
- Pastor L, Stambaugh RF. 2009. Predictive systems: Living with imperfect predictors. *Journal of Finance* **64**: 1583–1628.
- Pastor L, Stambaugh RF. 2012. Are stocks really less volatile in the long run? *Journal of Finance* **67**: 431–478.
- Pesaran MH, Timmermann AG. 2002. Market timing and return prediction under model instability. *Journal of Empirical Finance* **9**: 495–510.
- Pettenuzzo D, Timmermann AG. 2011. Predictability of stock returns and asset allocation under structural breaks. *Journal of Econometrics* **164**: 60–78.
- Shanken JA, Tamayo A. 2004. Dividend yield, risk, and mispricing: a Bayesian analysis. SSRN Working Paper 548364.
- Villani M. 2005. Inference in vector autoregressive models with an informative prior on the steady state. Sveriges Riksbank Working Paper 181.
- Villani M. 2009. Steady-state priors for vector autoregressions. *Journal of Applied Econometrics* **24**: 630–650.
- Wachter JA, Warusawitharana M. 2009. Predictable returns and asset allocation: should a skeptical investor time the market? *Journal of Econometrics* **148**: 162–178.
- Welch I. 2001. The equity premium consensus forecast revisited. Cowles Foundation Discussion Paper 1325.
- Xia Y. 2001. Learning about predictability: the effect of parameter uncertainty on dynamic asset allocation. *Journal of Finance* **56**: 205–246.
- Yu B, Mykland P. 1998. Looking at Markov samplers through cusum path plots: a simple diagnostic idea. *Statistics and Computing* **8**: 275–286.

APPENDIX A: TERM STRUCTURE OF RISK

For the first-order VAR in equation (3) conditional expectations follow directly as $E_t[y_{t+j}] = \mu + \mathbf{B}^j(y_t - \mu)$ and give the decomposition

$$y_{t+j} - E_t[y_{t+j}] = \sum_{i=0}^j \mathbf{B}^{j-i} \varepsilon_{t+i} \tag{A1}$$

Define cumulative variables $Y_{t+k} = \sum_{j=1}^k y_{t+j}$ and use equation (A1) to obtain

$$Y_{t+k} - E_t[Y_{t+k}] \equiv \sum_{j=1}^k (y_{t+j} - E_t[y_{t+j}]) = \sum_{i=1}^k \left(\sum_{j=i}^k \mathbf{B}^{j-i} \right) \varepsilon_{t+i} \tag{A2}$$

The conditional covariance matrix of Y_{t+k} is

$$\begin{aligned} \text{var}[Y_{t+k}] &= \sum_{i=1}^k \left(\sum_{j=i}^k \mathbf{B}^{j-i} \right) \Sigma \left(\sum_{j=i}^k \mathbf{B}^{j-i} \right)' \\ &= \sum_{\ell=1}^k \left(\sum_{j=0}^{\ell-1} \mathbf{B}^j \right) \Sigma \left(\sum_{j=0}^{\ell-1} \mathbf{B}^j \right)' \end{aligned} \tag{A3}$$

By selecting the relevant linear combinations of elements from the covariance matrix and dividing by k , we obtain the term structure of risk for real asset returns.

APPENDIX B: POSTERIOR DENSITIES

For the VAR (3) with the prior (22) on μ the simulation for $\theta = (\Sigma, \mathbf{B}, \mu)$ proceeds in a three-block Gibbs sampler as in Villani (2009):

1. The conditional posterior of Σ is inverted Wishart $iW(\hat{S}, T)$ with $\hat{S} = \sum_t e_t e_t'$ the residual sum of squares matrix conditional on μ and \mathbf{B} .
2. Since the prior on \mathbf{B} is flat, conditional on being in the stationary region, the conditional posterior is the matrix-variate normal $N(\hat{\mathbf{B}}, \Sigma \otimes \mathbf{X}_\mu^{-1})$, where $\hat{\mathbf{B}}$ are the OLS estimates conditional on μ , and the matrix $\mathbf{X}_\mu = \sum_t (y_{t-1} - \mu)(y_{t-1} - \mu)'$. We draw until we obtain a valid \mathbf{B} in the stationary region.
3. For the posterior of μ conditional on (\mathbf{B}, Σ) , define $\mathbf{A} = \mathbf{I} - \mathbf{B}$ and $h_t = y_t - \mathbf{B}y_{t-1}$. Then the conditional posterior follows as $\mu \sim N(\hat{\mu}, \mathbf{V})$ with conditional moments

$$\mathbf{V} = \left(T\mathbf{A}'\Sigma^{-1}\mathbf{A} + \kappa\Omega_0^{-1} \right)^{-1} \tag{B1}$$

$$\hat{\mu} = \mathbf{V} \left(\mathbf{A}'\Sigma^{-1}\bar{h} + \kappa\Omega_0^{-1}\mu_0 \right) \tag{B2}$$

and \bar{h} the sample average of h_t conditional on \mathbf{A} .

We simulate from the joint posterior by iterating over the sequence of conditional posteriors. For the initialization of the Gibbs sampler we use the OLS estimates. We start the Gibbs sampler using 2500 draws that we discard. For the risk and portfolio analysis we use the subsequent sample of 20,000 parameter draws. Conditional on each draw for the parameters we simulate two antithetic scenarios of future returns. In this way we create 40,000 scenarios of future returns.

The burn-in phase is chosen by visual inspection of the posterior draws and supported by the convergence tool of Yu and Mykland (1998). We use a standardized version of their cumsum statistic as suggested by Bauwens *et al.* (2003). For all priors the plot of the standardized version of the cumsum statistic converges smoothly and quickly to zero, especially after the burn-in phase, which indicates the convergence of the Monte Carlo chain.