

The Term Structure of the Risk-Return Tradeoff

John Y. Campbell and Luis M. Viceira¹

An abbreviated version of this paper was published in the *Financial Analysts Journal*, Volume 61, No. 1, January/February 2005. This paper has been awarded the 2005 Graham and Dodd Award for Excellence in Financial Writing, Financial Analysts Journal and CFA Institute.

¹Campbell: Department of Economics, Littauer Center 213, Harvard University, Cambridge MA 02138, USA, and NBER. Tel 617-496-6448, email john_campbell@harvard.edu. Viceira: Harvard Business School, Morgan Hall 395, Boston MA 02163. Tel 617-495-6331, email lviceira@hbs.edu. This material is based upon work supported by the National Science Foundation under Grant No. 0214061 to Campbell. Viceira acknowledges the financial support of the Division of Research of the Harvard Business School.

The Term Structure of the Risk-Return Tradeoff

Abstract

Recent research in empirical finance has documented that expected excess returns on bonds and stocks, real interest rates, and risk shift over time in predictable ways. Furthermore, these shifts tend to persist over long periods of time. In this paper we propose an empirical model that is able to capture these complex dynamics, yet is simple to apply in practice, and we explore its implications for asset allocation. Changes in investment opportunities can alter the risk-return tradeoff of bonds, stocks, and cash across investment horizons, thus creating a “term structure of the risk-return tradeoff.” We show how to extract this term structure from our parsimonious model of return dynamics, and illustrate our approach using data from the U.S. stock and bond markets. We find that asset return predictability has important effects on the variance and correlation structure of returns on stocks, bonds and T-bills across investment horizons.

JEL classification: G12.

1 Introduction

Recent research in empirical finance has documented that expected excess returns on bonds and stocks, real interest rates, and risk shift over time in predictable ways. Furthermore, these shifts tend to persist over long periods of time. Starting at least with the pioneering work of Samuelson (1969) and Merton (1969, 1971, 1973) on portfolio choice, financial economists have argued that asset return predictability can introduce a wedge between the asset allocation strategies of short- and long-term investors.

One important implication of time variation in expected returns is that investors, particularly aggressive investors, may want to engage in market-timing (or tactical asset allocation) strategies aimed at maximizing short-term return, based on the predictions of their return forecasting model. However, there is considerable uncertainty about the degree of asset return predictability and this makes it hard to identify the optimal market-timing strategy. A second, less obvious implication of asset return predictability is that risk—defined as the conditional variances and covariances per period of asset returns—may be significantly different across investment horizons, thus creating a “term structure of the risk-return tradeoff.” This tradeoff is the focus of this paper.

This paper examines the implications for risk across investment horizons of time variation in investment opportunities. To this end we propose an empirical model that is able to capture the complex dynamics of expected returns and risk, yet is simple to apply in practice. Specifically, we model interest rates and returns as a vector autoregressive model (VAR).² We show how one can easily extract the term structure of risk using this parsimonious model of return dynamics, and illustrate our approach using quarterly data from the U.S. stock, bond and T-bill markets for the postwar period. In our empirical application we use variables that have been identified as return predictors by past empirical research, such as the short-term interest rate, the dividend-price ratio, and the yield spread between long-term and short-term bonds.³

²This type of specification has been used in a similar context by Kandel and Stambaugh (1987), Campbell (1991), Hodrick (1992), Campbell and Viceira (1999), Barberis (2000), and Campbell, Chan and Viceira (2003) among others.

³See Fama and Schwert, (1977), Campbell (1987), Glosten, Jagannathan, and Runkle (1993) for evidence of predictability from the short-term interest rate; Campbell and Shiller (1988) and Fama and French (1988) for the dividend-price ratio; and Shiller, Campbell, and Schoenholtz (1983), Fama (1984), Fama and French (1989), and Campbell and Shiller (1991) for the yield spread.

These variables enable us to capture horizon effects on stock market risk, inflation risk, and real interest rate risk.

Campbell and Viceira (2002) have shown that empirically asset return predictability has important effects on the variances of long-horizon returns on stocks, bonds, and T-bills. Building on their work, we explore the correlation structure of asset returns across investment horizons. Correlations are just as important as variances for long-horizon investors.

We use mean-variance analysis to highlight the relevance of risk horizon effects on asset allocation. Traditional mean-variance analysis typically focuses on short-term expected returns and risk. We extend this analysis to a multi-horizon setting. In the context of our model we are able to show the limitations of traditional mean-variance analysis: It is valid only when the term structure of the risk-return tradeoff is flat; otherwise, it describes only the short end of this curve. We use our model to characterize the efficient mean-variance frontier at different investment horizons, by looking at the risk and composition of the global minimum-variance portfolio and a tangency portfolio of bonds and stocks.

In order to concentrate on risk horizon effects, we abstract from several other considerations that may be important in practice. We ignore changes in volatility through time; such changes are typically short-lived and have only a secondary influence on long-term risks. We consider only the first two moments of returns, ignoring the possibility that investors care also about other properties of the return distribution. And for simplicity, we show unconditional average portfolio allocations rather than the full range of allocations that would be optimal under different market conditions.

The concept of a term structure of the risk-return tradeoff is conceptually appealing but, strictly speaking, is only valid for buy-and-hold investors who make a one-time asset allocation decision and are interested only in the assets available for spending at the end of a particular horizon. In practice, however, few investors can truly be characterized as buy-and-hold. Most investors, both individuals and institutions such as pension funds and endowments, can rebalance their portfolios and have intermediate spending needs. An important open question is to what extent the simple strategies implied by the term structure of the risk-return tradeoff are a good approximation to more complex rebalancing strategies.

The organization of the paper is as follows. Section 2 introduces our dynamic

model of asset returns, and explores a simple example in which all moments of interest can be written out explicitly. Section 3 considers an empirical application with three U.S. asset classes and three return forecasting variables. Section 4 explores the implications of the model for risk across investment horizons. Section 5 extends traditional short-term mean-variance analysis to a multi-horizon setting, and Section 6 concludes. Details are given in a technical appendix, “Long-Horizon Mean-Variance Analysis: A User Guide.”

2 Risk and Return in a Vector Autoregressive Model

2.1 Dynamics of returns and state variables

We first describe a stylized model that can capture the expected return and risk of asset returns at different horizons. This model requires us to specify the asset classes under consideration, plus any variables that can help us form expectations of future returns on these asset classes—such as price-earnings ratios, interest rates, or yield spreads. We will refer to these variables as “state variables”. Our approach allows for any number of asset classes and state variables.

Let \mathbf{z}_{t+1} denote a column vector whose elements are the returns on all asset classes under consideration, and the values of the state variables at time $(t + 1)$. Because it is convenient for our subsequent portfolio analysis, we choose to write this vector as

$$\mathbf{z}_{t+1} \equiv \begin{bmatrix} r_{0,t+1} \\ \mathbf{r}_{t+1} - r_{0,t+1}\mathbf{1} \\ \mathbf{s}_{t+1} \end{bmatrix} \equiv \begin{bmatrix} r_{0,t+1} \\ \mathbf{x}_{t+1} \\ \mathbf{s}_{t+1} \end{bmatrix}. \quad (1)$$

where $r_{0,t+1}$ denotes the log (or continuously compounded) real return on the asset that we use as a benchmark to compute excess returns on all other asset classes, \mathbf{x}_{t+1} is a vector of log excess returns on all other asset classes with respect to the benchmark, and \mathbf{s}_{t+1} is a vector with the realizations of the state variables. For future reference, we assume there are $n + 1$ asset classes, and $m - n - 1$ state variables, so that \mathbf{z}_{t+1} has $(m \times 1)$ elements.

Note that all the returns included in \mathbf{z}_{t+1} are continuously compounded (or log) returns instead of gross returns. We work with log returns because it is more convenient from a data-modeling perspective. Of course, investors are concerned about

gross returns rather than log returns. Thus in our portfolio analysis we reverse the log transformation whenever it is necessary.

The choice of a benchmark asset is arbitrary; we normally choose the benchmark to be a Treasury bill since this is the asset with the smallest short-term risk, but the representation of returns in (1) is perfectly general. We could just as easily write the vector in terms of real returns rather than excess returns. We show how to extract the moments of real returns from this VAR in the appendix.

Our key assumption about the dynamics of asset returns and state variables is that they follow a first-order vector autoregressive process, or VAR(1). Each variable $z_{i,t+1}$ included in \mathbf{z}_{t+1} depends linearly on a constant, its own lagged value, the lagged value of all other variables in \mathbf{z}_{t+1} , and a contemporaneous random shock $v_{i,t+1}$:

$$z_{i,t+1} = \phi_0 + \phi_1 z_{1,t} + \dots + \phi_i z_{i,t} + \dots + v_{i,t+1}. \quad (2)$$

Note that this is simply a generalization of a first-order autoregressive process—or AR(1)—to handle multiple forecasting variables. In an AR(1), all coefficients in equation (2) are zero, except ϕ_0 and ϕ_i .

Stacking together all these forecasting equations, we can represent the VAR(1) compactly as

$$\mathbf{z}_{t+1} = \mathbf{\Phi}_0 + \mathbf{\Phi}_1 \mathbf{z}_t + \mathbf{v}_{t+1}, \quad (3)$$

where $\mathbf{\Phi}_0$ is a vector of intercepts; $\mathbf{\Phi}_1$ is a square matrix that stacks together the slope coefficients; and \mathbf{v}_{t+1} is a vector of zero-mean shocks to realizations of returns and return forecasting variables.

We assume that the matrix of slopes $\mathbf{\Phi}_1$ is well behaved in a statistical sense, by requiring that its determinant is bounded between -1 and $+1$. This is the multivariate equivalent of the stationarity condition in a AR(1), that requires the autoregressive parameter to be bounded between -1 and $+1$. This condition ensures that, in the absence of shocks, the variables that enter the VAR(1) converge to their long-run means in a finite number of periods. Thus this condition excludes explosive (or nonstationary) behavior in these variables.

Finally, to complete the description of the return dynamics we need to be more specific about the nature of the vector of shocks to asset returns and return forecasting variables (\mathbf{v}_{t+1}). In particular, we assume that the vector of shocks is normally distributed,

$$\mathbf{v}_{t+1} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \mathbf{\Sigma}_v), \quad (4)$$

where Σ_v denotes the matrix of contemporaneous variances and covariances of shocks. This matrix is not necessarily diagonal: We allow unexpected realizations of excess returns on different asset classes to covary with each other, and with shocks to return forecasting variables. For example, it is plausible that the excess return on domestic equity is correlated with the excess return on foreign equity, or that unexpected shocks to interest rates (a possible state variable) are correlated with domestic equity returns.

For future reference, we note here that, consistent with our representation of \mathbf{z}_{t+1} in (1), we can write Σ_v as

$$\Sigma_v \equiv \text{Var}_t(\mathbf{v}_{t+1}) = \begin{bmatrix} \sigma_0^2 & \sigma'_{0x} & \sigma'_{0s} \\ \sigma_{0x} & \Sigma_{xx} & \Sigma'_{xs} \\ \sigma_{0s} & \Sigma_{xs} & \Sigma_s \end{bmatrix},$$

where the elements on the main diagonal are the variance of the real return on the benchmark asset (σ_0^2), the variance-covariance matrix of unexpected excess returns (Σ_{xx}), and the variance-covariance matrix of the state variables (Σ_s). The off diagonal elements are the covariances of the real return on the benchmark assets with excess returns on all other assets and with shocks to the state variables (σ_{0x} and σ_{0s}), and the covariances of excess returns with shocks to the state variables (Σ_{xs}).

To keep our exposition as simple as possible, we assume that these variances and covariances do not vary over time—in other words, that risk does not change over time. However, we want to emphasize here that this assumption, while perhaps not realistic, is nevertheless not constraining from the perspective of long-term portfolio choice. The empirical evidence available suggests that changes in risk are a short-lived phenomenon.⁴ In this sense, time varying risk is unlikely to be a major concern to long-term investors.

It is useful to note that the unconditional mean vector and variance-covariance

⁴Authors such as Campbell (1987), Harvey (1989, 1991), and Glosten, Jagannathan, and Runkle (1993) have explored the ability of the state variables used here to predict risk and have found only modest effects that seem to be dominated by the effects of the state variables on expected returns. Chacko and Viceira (1999) show how to include changing risk in a long-term portfolio choice problem, using a continuous-time extension of the methodology of Campbell and Viceira (1999); they find that changes in equity risk are not persistent enough to have large effects on the intertemporal hedging demand for equities. Aït-Sahalia and Brandt (2001) adopt a semiparametric methodology that accommodates both changing expected returns and changing risk.

matrix of \mathbf{z}_{t+1} are given by

$$\begin{aligned}\boldsymbol{\mu}_z &= (\mathbf{I}_m - \boldsymbol{\Phi}_1)^{-1} \boldsymbol{\Phi}_0, \\ \text{vec}(\boldsymbol{\Sigma}_{zz}) &= (\mathbf{I}_{m^2} - \boldsymbol{\Phi}_1 \otimes \boldsymbol{\Phi}_1)^{-1} \text{vec}(\boldsymbol{\Sigma}_v).\end{aligned}\tag{5}$$

These equations clarify why we require that the determinant of the matrix of slopes $\boldsymbol{\Phi}_1$ is bounded between -1 and 1 . Otherwise the unconditional mean and variance of \mathbf{z}_{t+1} would not be defined.

2.2 The VAR(1) model of return dynamics and the traditional view of the risk-return tradeoff

It is instructive to contrast the VAR(1) model with the traditional view that assets have constant expected returns, which can be estimated from their historical means. The traditional approach is a restricted version of the VAR(1) model in which only a constant term is used to forecast returns.

An investor who uses the VAR(1) model to forecast returns will perceive the risk-return tradeoff differently than an investor who uses the traditional approach. First, the VAR(1) investor will have a different return expectation each period, based on the changing state variables of the model. This return expectation is known as a conditional expectation to distinguish it from the unconditional expectation used in the traditional approach.

Second, the VAR(1) investor will measure the short-term risk of each asset by its variance relative to its conditional expectation, rather than its unconditional expectation. Even if this conditional variance is constant over time, as we assume in our empirical work, it differs from the unconditional variance. The VAR(1) investor understands that some portion of the unconditional volatility of each asset return is actually predictable time-variation in the return and thus does not count as risk. For this reason the conditional variance is smaller than the unconditional variance.

Third, the VAR(1) investor will understand that the long-term risks of asset returns may differ from their short-term risks. In the traditional approach, with constant expected returns, the variance of each asset return is proportional to the horizon over which it is held. A one-year variance is four times a one-quarter variance, and a decadal variance is ten times larger again. Annualized variances are independent

of the time horizon, and thus there is a single number that summarizes risks for all holding periods. In the VAR(1) model, by contrast, annualized variances may either increase or decline as the holding period increases. Annualized covariances may also be sensitive to the holding period, and these effects of the investment horizon on perceived risk are the main subject of this paper.

To illustrate these effects in an explicit but reasonably simple way, we now present a special case of the VAR(1) model in which there is a single state variable that predicts two different asset returns. The best forecasts of the state variable and the asset returns are all proportional to the state variable, so the lagged asset returns do not contribute forecasting power as they would in the general VAR(1) model. In this special case the formulas for conditional and unconditional expectations, conditional and unconditional short-term variances, and longer-term variances, can all be written out so one can see how they differ.

2.3 Risk and return in a VAR(1) model: A simple example

The appendix presents detailed derivations of the formulas for conditional expectations, variances and covariances of returns across investment horizons when the dynamics of asset returns and state variables are described by a general VAR(1) model. These formulas are somewhat complicated, but we can illustrate the effects of the investment horizon on risk by restricting attention to a special case of the VAR(1) model in which there is a single state variable that predicts two different asset returns (equities and bonds).

In our example, the dynamics of returns and the state variable are given by

$$\begin{aligned}
 re_{t+1} &= \phi_{es}s_t + v_{e,t+1}, \\
 rb_{t+1} &= \phi_{bs}s_t + v_{b,t+1}, \\
 s_{t+1} &= \phi_{ss}s_t + v_{s,t+1},
 \end{aligned} \tag{6}$$

where re_{t+1} denotes the log return on equities, rb_{t+1} denotes the log return on bonds, and s_{t+1} denotes the state variable. For stationarity, the autoregressive coefficient ϕ_{ss} must be less than one in absolute value.

In the model (6), the predictable component of returns is given by the terms $\phi_{es}s_t$ and $\phi_{bs}s_t$. These terms describe the expected return on equities and bond one

period ahead. The unpredictable or unexpected component of realized returns and the return forecasting variable is given by $v_{e,t+1}$, $v_{b,t+1}$, and $v_{s,t+1}$. That is, $v_{e,t+1} = re_{t+1} - \text{E}_t[re_{t+1}]$ and likewise $v_{b,t+1} = rb_{t+1} - \text{E}_t[rb_{t+1}]$, and $v_{s,t+1} = s_{t+1} - \text{E}_t[s_{t+1}]$.

To complete the description of return dynamics we need to be more specific about the nature of the shocks $v_{e,t+1}$, $v_{b,t+1}$, and $v_{s,t+1}$. We assume that these shocks are serially uncorrelated but have a contemporaneous variance-covariance matrix

$$\text{Var}_t \begin{pmatrix} v_{e,t+1} \\ v_{b,t+1} \\ v_{s,t+1} \end{pmatrix} = \begin{bmatrix} \sigma_e^2 & \sigma_{eb} & \sigma_{es} \\ \sigma_{eb} & \sigma_b^2 & \sigma_{bs} \\ \sigma_{es} & \sigma_{bs} & \sigma_s^2 \end{bmatrix} \equiv \Sigma_v. \quad (7)$$

We now examine how asset return predictability generates horizon effects in risk, in the context of the simple example (6). In the spirit of traditional one-period mean-variance analysis, we define risk per period at horizon k as the conditional variance-covariance matrix of k -period log asset returns, divided by the length of the horizon:

$$\frac{1}{k} \text{Var}_t \begin{pmatrix} re_{t+k}^{(k)} \\ rb_{t+k}^{(k)} \end{pmatrix} = \frac{1}{k} \begin{bmatrix} \text{Var}_t \left(re_{t+k}^{(k)} \right) & \text{Cov}_t \left(re_{t+k}^{(k)}, rb_{t+k}^{(k)} \right) \\ \text{Cov}_t \left(re_{t+k}^{(k)}, rb_{t+k}^{(k)} \right) & \text{Var}_t \left(rb_{t+k}^{(k)} \right) \end{bmatrix}. \quad (8)$$

Here $re_{t+k}^{(k)}$ and $rb_{t+k}^{(k)}$ denote the k -horizon log return on equities and bonds, respectively. It is useful to note for future reference that k -period log returns are sums of one-period log returns over k successive periods. For example,

$$re_{t+k}^{(k)} = re_{t+1} + \dots + re_{t+k}, \quad (9)$$

where $re_{t+1} \equiv re_{t+1}^{(1)}$ is the one-period log return on equities.

When $k = 1$, equation (8) reduces to the variance-covariance matrix of one period returns, which is the standard measure of portfolio risk in mean-variance analysis. This variance-covariance matrix is given by:

$$\text{Var}_t \begin{pmatrix} re_{t+1} \\ rb_{t+1} \end{pmatrix} = \text{Var}_t \begin{pmatrix} v_{e,t+1} \\ v_{b,t+1} \end{pmatrix} = \begin{bmatrix} \sigma_e^2 & \sigma_{eb} \\ \sigma_{eb} & \sigma_b^2 \end{bmatrix}. \quad (10)$$

The first equality in equation (10) follows from the fact that a conditional variance is computed over deviations of the variable of interest with respect to its conditional mean; for next-period equity returns, this deviation is

$$re_{t+1} - \text{E}_t[re_{t+1}] = v_{e,t+1}, \quad (11)$$

and likewise for rb_{t+1} . The second equality in equation (10) follows directly from equation (7).

When the investment horizon goes beyond one period, (8) does not generally reduce to (10) except in the special but important case where returns are not predictable. This in turn implies that the risk structure of returns varies across investment horizons. To illustrate this, we explore the variances and covariances of two-period returns in detail, and note that the results for this case extend to longer horizons.

We start with the variance of two-period equity returns. Using (9), we can decompose this variance in terms of the variances and autocovariances of one-period equity returns as follows:

$$\begin{aligned} \frac{1}{2} \text{Var}_t \left(re_{t+2}^{(2)} \right) &= \frac{1}{2} \text{Var}_t (re_{t+1} + re_{t+2}) \\ &= \frac{1}{2} \text{Var}_t (re_{t+1}) + \frac{1}{2} \text{Var}_t (re_{t+2}) + \text{Cov}_t (re_{t+1}, re_{t+2}) \end{aligned} \quad (12)$$

The decomposition of the variance of two-period bond returns is identical, except that we replace re with rb . Thus the variance per period of two-period returns depends on the variance of one-period returns next period and two periods ahead, and on the serial covariance of returns.

Direct examination of equation (12) for the two-period case reveals under which conditions there are no horizon effects in risk: When the variance of single period returns is the same at all forecasting horizons, and when returns are not autocorrelated. This guarantees that the variance per period of k -period returns remains the same for all horizons.

It turns out these conditions do not hold when returns are predictable, so horizon matters for risk. This is easy to see in the context of our simplified example (6). For this model we have that

$$re_{t+2} - \text{E}_t[re_{t+2}] = v_{e,t+2} + \phi_{es}v_{s,t+1}. \quad (13)$$

A similar expression obtains for bond returns. Equation (13) says that, from today's perspective, future unexpected returns depend not only on their own shocks, but also on lagged shocks to the forecasting variable.

Plugging (11) and (13) into (12) and computing the moments we obtain an explicit

expression for (12):

$$\frac{1}{2} \text{Var}_t \left(re_{t+2}^{(2)} \right) = \frac{1}{2} \sigma_e^2 + \frac{1}{2} (\sigma_e^2 + \phi_{es}^2 \sigma_s^2) + \phi_{es} \sigma_{es}. \quad (14)$$

In general, this expression does not reduce to $\text{Var}_t (re_{t+1}) = \sigma_e^2$ unless returns are not predictable—i.e. unless $\phi_{es} = 0$.

The terms in equation (14) correspond one for one to those in equation (12). They show that return predictability has two effects on the variance of multiperiod returns. First, it increases the conditional variance of future single period returns, relative to the variance of next period returns, because future returns depend on past shocks to the forecasting variable. The second term of the expression captures this effect.

Second, it induces autocorrelation in single period returns, because future single period returns depend on past shocks to the forecasting variable, which in turn are contemporaneously correlated with returns. The third term of the equation captures this effect. The sign of this autocorrelation is equal to the sign of the product $\phi_{es} \sigma_{es}$. For example, when s_t forecasts returns positively ($\phi_{es} > 0$), and the contemporaneous correlation between unexpected returns and the shocks to s_{t+1} is negative (σ_{es}), return predictability generated by s_t induces negative first-order autocovariance in one period returns. This is known as “mean-reversion” in returns.

These effects are present at any horizon. Equation (14) generalizes to the k -horizon case in the sense that the expression for $\text{Var}_t(re_{t+k}^{(k)})/k$ is equal to σ_e^2 plus two terms: a term in σ_s^2 which is always positive, and a term in $\phi_{es} \sigma_{es}$ which may be positive or negative, depending on the sign of this product. It is possible to show that the marginal contribution of these two terms to the overall variance per period increases at rates proportional to ϕ_{ss}^{2k}/k and ϕ_{ss}^k/k , respectively. Since $|\phi_{ss}| < 1$, this marginal contribution becomes negligible at long horizons, with the contribution of the term in σ_s^2 becoming negligible sooner. However it is important to note that the closer is $|\phi_{ss}|$ to one, that is, the more persistent is the forecasting variable, the longer it takes for its marginal contribution to the variance to disappear. Taken together, all of this implies that the variance per period of multiperiod returns, when plotted as a function of the horizon k , may be increasing, decreasing or hump-shaped at intermediate horizons, and it eventually converges to a constant—the unconditional variance of returns—at long enough horizons.

Similarly, we can decompose the covariance per period of two-period bond and

equity returns as

$$\begin{aligned}
\frac{1}{2} \text{Cov}_t \left(re_{t+2}^{(2)}, rb_{t+2}^{(2)} \right) &= \frac{1}{2} \text{Cov}_t (re_{t+1} + re_{t+2}, rb_{t+1} + rb_{t+2}) \\
&= \frac{1}{2} \text{Cov}_t (re_{t+1}, rb_{t+1}) + \frac{1}{2} \text{Cov}_t (re_{t+2}, rb_{t+2}) \quad (15) \\
&\quad + \frac{1}{2} [\text{Cov}_t (re_{t+2}, rb_{t+1}) + \text{Cov}_t (re_{t+1}, rb_{t+2})],
\end{aligned}$$

which, for the simple example (6), specializes to

$$\frac{1}{2} \text{Cov}_t \left(re_{t+2}^{(2)}, rb_{t+2}^{(2)} \right) = \frac{1}{2} \sigma_{eb} + \frac{1}{2} (\sigma_{eb} + \phi_{es} \phi_{bs} \sigma_s^2) + \frac{1}{2} (\phi_{es} \sigma_{bs} + \phi_{bs} \sigma_{es}). \quad (16)$$

Equation (16) obtains after plugging (11) and (13) into (15) and computing the moments. Once again, we find that the covariance per period of long-horizon returns does not equal the covariance of one-period returns (σ_{eb}) unless returns are unpredictable—i.e., unless $\phi_{es} = 0$ and $\phi_{bs} = 0$.

Asset return predictability has two effects on the covariance of multiperiod returns. These effects are similar to those operating on the variance of multiperiod returns. The first one is the effect of past shocks to the forecasting variable on the contemporaneous covariance of future one-period stock and bond returns. This corresponds to the second term in (15) and (16). As equation (13) shows, a shock to s_{t+1} feeds to future stock and bond returns through the dependence of stock and bond returns on s_{t+1} . If s_{t+1} forecasts future stock and bond returns with the same sign (i.e., $\phi_{es} \phi_{bs} > 0$), a shock to s_{t+1} implies that bond and stock returns will move in the same direction in future periods; if s_{t+1} forecasts future stock and bond returns with different signs, a shock to s_{t+1} implies that bond and stock returns will move in opposite directions in future periods. Thus, asset return predictability implies that $\text{Cov}_t (re_{t+2}, rb_{t+2})$ may be larger or smaller than $\text{Cov}_t (re_{t+1}, rb_{t+1})$, depending on the sign and magnitude of $\phi_{es} \phi_{bs} \sigma_s^2$.

Asset return predictability also induces cross autocorrelation in returns. Future stock (bond) returns are correlated with lagged bond (stock) returns through their dependence on past shocks to the forecasting variable, and their contemporaneous correlation with shocks to this variable. This effect is captured by the third term in equations (15) and (16). For example, suppose that s_{t+1} forecasts negatively stock returns (i.e., $\phi_{es} < 0$) and that realized bond returns and shocks to s_{t+1} are negatively correlated ($\sigma_{bs} < 0$). Then equity returns two period ahead will be positively

correlated with bond returns one period ahead: A positive bond return one period ahead will tend to coincide with a negative shock to s_{t+1} , which in turn forecasts a positive stock return two periods ahead.

These effects are present at any horizon, and can lead to different shapes of the covariance of multiperiod returns at intermediate horizons, as one effect reinforces the other, or offsets it. Of course, at long enough horizons, the covariance per period eventually converges to a constant—the unconditional covariance of returns.

2.4 Extending the basic VAR(1) model

There are several important extensions of the VAR(1) model described in Section 2.1. First, it is straightforward to allow for any number of lags of asset returns and forecasting variables within the VAR framework. It is interesting however to note that one can rewrite a VAR of any order in the form of a VAR(1) by adding more state variables which are simply lagged values of the original vector of variables. Thus all the formal results in this paper and in the companion technical document are valid for any VAR specification, even if for convenience we write them only in terms of a VAR(1).

An important consideration when considering additional lags is whether the parameters of the VAR are unknown and must be estimated from historical data. In that case, the precision of those estimates worsens as the number of parameters increases relative to the size of the sample. Since the number of parameters in the model increases exponentially with the number of lags, adding lags can reduce significantly the precision of the parameter estimates.⁵ Of course, similar concerns arise with respect to the number of asset classes and state variables included in the VAR.

Second, some variables commonly used to forecast asset returns are highly persistent and have innovations that are highly correlated with stock returns. This implies that OLS estimates of the VAR parameters may be biased in finite samples (Stambaugh 1999). A standard econometric procedure to correct for these biases is bootstrapping. Bias-corrected estimates typically show that there is less predictabil-

⁵For example, if there are three asset classes and three state variables (as we are doing in our empirical application), the VAR(1) specification requires estimating a total of 63 parameters—6 intercepts, 36 slope coefficients, 6 variances, and 15 covariance terms. Adding an additional lag for all variables would increase this already fairly large number of parameters to 99 (a 57% increase).

ity of excess stock returns than OLS estimates suggest (Hodrick, 1992; Goetzmann and Jorion, 1993; Nelson and Kim, 1993), but greater predictability of excess bond returns (Bekaert, Hodrick, and Marshall, 1997). The reason for the discrepancy is that the evidence on stock market predictability comes from positive regression coefficients of stock returns on the dividend-price ratio, while the evidence on bond market predictability comes from positive regression coefficients of bond returns on yield spreads. Stambaugh (1999) shows that the small-sample bias in such regressions has the opposite sign to the sign of the correlation between innovations in returns and innovations in the predictive variable. In the stock market the log dividend-price ratio is negatively correlated with returns, leading to a positive small-sample bias which helps to explain some apparent predictability; in the bond market, on the other hand, the yield spread is mildly positively correlated with returns, leading to a negative small-sample bias which cannot explain the positive regression coefficient found in the data.

It is important to note that bootstrapping methods are not a panacea when forecasting variables are persistent, because we do not know the true persistence of the data generating process that is to be used for the bootstrap (Elliott and Stock, 1994, Cavanagh, Elliott, and Stock, 1995, Campbell and Yogo, 2003). Although finite-sample bias may well have some effect on the coefficients reported in our empirical analysis below, we do not attempt any bias corrections in this paper. Instead, we take the estimated VAR coefficients as given and known by investors, and explore their implications for optimal long-term portfolios.

Third, we have assumed that the variance-covariance structure of the shocks in the VAR is constant. However, one can allow for time variation in risk. We have argued that this might not be too important a consideration from the perspective of long-term asset allocation, because the empirical evidence available suggests that changes in risk exhibit very low persistence. Nonetheless, time-varying risk can be incorporated into the model, for example, along the lines of the models suggested by Bollerslev (1990), Engle (2002), or Rigobon and Sack (2003).

Fourth, one might have prior views about some of the parameters of the VAR model, particularly about the intercepts which determine the long-run expected return on each asset class, or the slope coefficients which determine the movements in expected returns over time. Using Bayesian methods one can introduce these prior views and combine them with estimates from the data. Hoevenaars, Molenaar, and Steenkamp (2004) examine this extension of the model in great detail.

3 Return Dynamics of U.S. Bills, Bonds and Equities

To illustrate our approach to the modelling of asset return dynamics, we consider a practical application with three U.S. asset classes and three return forecasting variables—in addition to the lagged returns on the three asset classes. The assets are cash, equities, and Treasury bonds. The return forecasting variables are the log short-term nominal interest rate, the log dividend yield, and the slope of the yield curve—or yield spread. Previous empirical research has shown that these variables have some power to forecast future excess returns on equities and bonds.

Because we do not know the parameters that govern the relation between these variables, and we do not want to impose any prior views on these parameters, we estimate the VAR using quarterly data from the Center for Research and Security Prices (CRSP) of the University of Chicago for the period 1952.Q2–2002.Q4. We choose 1952.Q2 as our starting date because it comes shortly after the Fed-Treasury Accord that allowed short-term nominal interest rates to freely fluctuate in the market. However, the main results shown here are robust to other sample periods that include the pre-WWII period, or that focus only on the last 25 years.

Following standard practice, we consider cash as our benchmark asset and compute log excess returns on equities and Treasury bonds with respect to the real return on cash. We proxy the real return on cash by the ex-post real return on 90-day T-bills (i.e., the difference between the log yield on T-bills and the log inflation rate). We also use the log yield on the 90-day T-bill as our measure of the log short-term nominal interest rates.⁶ The log return on equities (including dividends) and the log dividend yield are for a portfolio that includes all stocks traded in the NYSE, NASDAQ, and AMEX markets. The log return on Treasury bonds is the log return on a constant maturity 5-year Treasury bond, and the yield spread is the difference between the log yield on a zero-coupon 5-year Treasury bond and the yield on a 90-day T-bill. This is an updated version of the VAR estimated by Campbell, Chan and Viceira (2003).

Table 1 shows the sample mean and standard deviation of the variables included in the VAR. Except for the log dividend yield, the sample statistics are in annualized,

⁶Note that by including both the real and the nominal ex-post log short-term interest rate in the VAR we can also capture the dynamics of inflation, since log inflation is simply the difference between the log nominal interest rate and the ex-post log real interest rate.

percentage units. We adjust mean log returns by adding one-half their variance so that they reflect mean gross returns. For the post-war period, Treasury bills offer a low average real return (a mere 1.52% per year) along with low variability. Stocks have an excess return of 6.31% per year compared to 1.37% for the 5-year bond. Although stock return volatility is considerably higher than bond return volatility (16.92% vs. 5.65%), the Sharpe ratio is two and a half times as high for stocks as for bonds. The average Treasury bill rate and yield spread are 5.37% and 1.03%, respectively.

Our estimates of the VAR, shown in Table 2, update and confirm the findings in Campbell, Chan and Viceira (2003). Table 2 reports the estimation results for the VAR system. The top section of the table reports coefficient estimates (with t -statistics in parentheses) and the R^2 statistic for each equation in the system. We do not report the intercept of each equation because we estimate the VAR imposing the restriction that the unconditional means of the variables implied by the VAR coefficient estimates equal their full-sample arithmetic counterparts.⁷ The bottom section of each panel shows the covariance structure of the innovations in the VAR system. The entries above the main diagonal are correlation statistics, and the entries on the main diagonal are standard deviations multiplied by 100. All variables in the VAR are measured in natural units, so standard deviations are per quarter.

The first row of each panel corresponds to the real bill rate equation. The lagged real bill rate and the lagged nominal bill rate have positive coefficients and highly significant t -statistics. The yield spread also has a positive coefficient and a t -statistic above 2.0 in the quarterly data. Thus a steepening of the yield curve forecasts an increase in the short-term real interest rate next period. The remaining variables are not significant in predicting real bill rates one period ahead.

The second row corresponds to the equation for the excess stock return. The lagged nominal short-term interest rate (with a negative coefficient) and the dividend-price ratio (with a positive coefficient) are the only variables with t -statistics above 2.0. Predicting excess stock returns is difficult: this equation has the lowest R^2 (9.5%). It is important to emphasize though that this low quarterly R^2 can be misleading about the magnitude of predictability at lower frequencies (say, annual). Campbell (2001) notes that when return forecasting variables are highly persistent the implied annual R^2 can be several times the reported quarterly R^2 . As we note

⁷Standard, unconstrained least-squares fits exactly the mean of the variables in the VAR excluding the first observation. We use constrained least-squares to ensure that we fit the full-sample means.

below, this is the case for the dividend yield and the short rate, the two main stock return forecasting variables.

The third row is the equation for the excess bond return. The yield spread, with a positive coefficient, is the only variable with a t -statistic well above 2.0. Excess stock returns, with a negative coefficient, also help predict future excess bond returns, but the t -statistic is only marginally significant. The R^2 is only 9.7%, slightly larger than the R^2 of the excess stock return equation. Once again, the bond excess return forecasting variable is highly persistent, which implies that bond return predictability is likely to be much larger at lower frequencies.

The last three rows report the estimation results for the remaining state variables, each of which are fairly well described by a persistent univariate AR(1) process. The nominal bill rate in the fourth row is predicted by the lagged nominal yield, whose coefficient is above 0.9, implying extremely persistent dynamics. The log dividend-price ratio in the fifth row also has persistent dynamics; the lagged dividend-price ratio has a coefficient of 0.96. The yield spread in the sixth row also seems to follow an AR(1) process, but it is considerably less persistent than the other variables.

The bottom section of the table describes the covariance structure of the innovations in the VAR system. Unexpected log excess stock returns are positively correlated with unexpected log excess bond returns, though this correlation is fairly low. Unexpected log excess stock returns are highly negatively correlated with shocks to the log dividend-price ratio. Unexpected log excess bond returns are highly negatively correlated with shocks to the nominal bill rate, but positively correlated with shocks to the ex-post short-term real interest rate, and mildly positively correlated with shocks to the yield spread.

4 The Risk of Equities, Bonds, and Bills Across Investment Horizons

In this section we examine in detail the implications of asset return predictability for risk at different horizons, the term structure of the risk-return tradeoff. Figures 1 and 2 illustrate the effect of investment horizon on the annualized risks of equities, bonds and bills. These figures are based on the VAR(1) estimates shown in Table 2, from which we have computed the conditional variances and covariances per period of

real returns across investment horizons. Figure 1 plots percent annualized standard deviations (the square root of variance per quarter times 200) of real returns for investment horizons up to 50 years, and Figure 2 plots percent correlations. Note that we are not looking directly at the long-horizon properties of returns, but at the long-horizon properties of returns imputed from our first-order VAR. Thus, provided that our VAR captures adequately the dynamics of the data, we can consistently estimate the moments of returns over any desired horizon.

These figures plot standard deviations and correlations for real returns on T-bills, equities, constant maturity 5-year Treasury bonds, and a zero-coupon nominal Treasury bond with k years to maturity that is held to maturity. The only uncertainty about the k -year real return on this variable-maturity bond is inflation, since the nominal principal is guaranteed. Thus the unexpected log real k -year return on the variable-maturity bond is just the negative of unexpected cumulative inflation from time t to time $t + k$.

We have noted in Section 2.1 that, if returns are unpredictable, their risks per period and correlations are the same across investment horizons. Thus in that case we should see flat lines in Figures 1 and 2. Far from that, most lines in those figures have slopes that change with the investment horizon, reflecting the predictability of real returns implied by the VAR(1) system of Table 2.

Figure 1 shows that long-horizon returns on stocks are significantly less volatile than their short-horizon returns. This strong decline in volatility is the result of mean-reverting behavior in stock returns induced by the predictability of stock returns from the dividend yield: The large negative correlation of shocks to the dividend yield and unexpected stock returns, and the positive significant coefficient of the log dividend yield in the stock return forecasting equation imply that low dividend yields tend to coincide with high current stock returns, and forecast poor future stock returns. Mean-reversion in stock returns cuts the annualized standard deviation of returns from about 17% per annum to less than 8% as one moves from a one-quarter horizon to a 25-year horizon.

The return on the 5-year bond also exhibits slight mean-reversion, with volatility declining from about 6% per annum at a one-quarter horizon, to 4% per annum at long horizons. This slight mean-reversion is the result of two offsetting effects. On the one hand the yield spread forecasts bond returns positively, and its shocks exhibit low positive correlation with unexpected bond returns. This per se causes mean-reversion in bond returns. On the other hand the nominal T-bill yield forecasts excess bond

returns positively, and its shocks are highly negatively correlated with unexpected bond returns. This causes mean-reversion in bond returns. The effect of the nominal bill yield ultimately dominates because it exhibits much more persistence, and it is more volatile than the yield spread. Note however, that the coefficient on the nominal rate is not statistically significant, while the coefficient on the yield spread is highly significant.

In contrast to the mean-reversion displayed by the real returns on stocks and the constant-maturity bond, the real returns on both T-bills and the variable-maturity bond exhibit mean-aversion. That is, their real return volatility increases with the investment horizon. The mean-aversion of T-bill returns is caused by persistent variation in the real interest rate in the postwar period, which amplifies the volatility of returns when Treasury bills are reinvested over long horizons. Campbell and Viceira (2002) have noted that mean-aversion in T-bill returns is even more dramatic in the pre-war period, when T-bills actually become riskier than stocks at sufficiently long investment horizons, a point emphasized by Siegel (1994).

The increase in return volatility at long horizons is particularly large for the variable-maturity bond whose initial maturity is equal to the holding period. Since the risk of this bond is the risk of cumulative inflation over the investment horizon, this reflects significant persistent variation in inflation in the postwar period. A positive shock to inflation that lowers the real return on a long-term nominal bond is likely to be followed by high inflation in subsequent periods as well, and this amplifies the annualized volatility of a long-term nominal bond held to maturity. Thus inflation risk makes a strategy of buying and holding long-term nominal bonds riskier than holding shorter term nominal bonds at all horizons. At long horizons, this strategy is even riskier than holding stocks. At horizons of up to 30 years, stocks are still riskier than bills and bonds. However the relative magnitude of these risks changes with the investment horizon.

Figure 2 shows that the correlation structure of real returns also exhibits interesting patterns across investment horizons.⁸ Real returns on stocks and fixed-maturity bonds are positively correlated at all horizons, but the magnitude of their correlation changes dramatically across investment horizons. At short horizons of a few quarters, correlation is about 20%, but it quickly increases to 60% at horizons of about six

⁸Since correlation is the ratio of covariance to the product of standard deviations, patterns in correlations do not have to be the same as patterns in covariances. However, in this case they are, and we report correlations instead of covariances because of their more intuitive interpretation.

years, and stays above 40% for horizons up to 18 years; at longer horizons, it declines steadily to levels around 15%. Of course, it is difficult to put too much weight on the effects predicted by the model at very long horizons, because of the size of our sample, but the increasing correlation at the short and medium horizons is certainly striking. Results for raw real returns not shown here also exhibit an increasing correlation pattern at those horizons.

This striking pattern in the correlation of multiperiod returns on stocks and bonds is the result of the interaction of two state variables that dominate at different horizons. At intermediate horizons, the most important variable is the short-term nominal interest rate, the yield on T-bills. Table 2 shows that the T-bill yield moves in a fairly persistent fashion. It predicts low returns on stocks, and its movements are strongly negatively correlated with bond returns. When the T-bill yield increases, bond returns fall at once, while stock returns react more slowly. Thus the intermediate-term correlation between bonds and stocks is higher than the short-term correlation because it takes time for interest-rate changes to have their full effect on stock prices.

At long horizons, the most important variable is the dividend-price ratio because this is the most persistent variable in our empirical model. The dividend-price ratio predicts high returns on stocks and low returns on bonds. In the very long run this weakens the correlation between stock and bond prices, because decades with a high dividend-price ratio will tend to have high stock returns and low bond returns, while decades with a low dividend-price ratio will tend to have low stock returns and high bond returns.

These two state variables also generate an interesting pattern in the correlation of stock returns with nominal bonds held to maturity. Over very short periods, stocks are only weakly correlated with bonds held to maturity but the correlation rises to a maximum of 62% at a horizon of 7 years. It then falls again and eventually turns negative. Recall that the uncertainty in the returns on nominal bonds held to maturity is entirely due to uncertainty about cumulative inflation to the maturity date. Thus real stock returns are weakly negatively correlated with inflation at short horizons, strongly negatively correlated at intermediate horizons, and weakly positively correlated with inflation at very long horizons. This is consistent with evidence that inflation creates stock market mispricing that can have large effects at intermediate horizons, but eventually corrects itself (Modigliani and Cohn, 1979, Ritter and Warr, 2002, Campbell and Vuolteenaho, 2004). In the very long run stocks are real assets and are able to hedge inflation risk.

5 Mean-Variance Allocations Across Investment Horizons

We have shown in Section 4 that asset return predictability can have dramatic effects on the variances and covariances per period of asset returns across investment horizons. These results are directly relevant for buy-and-hold investors with fixed investment horizons. In this section we use mean-variance analysis to highlight this relevance since, at any given horizon, the mean-variance efficient frontier is the set of buy-and-hold portfolios with minimum risk (or variance) per expected return. Throughout this section we will consider the set of efficient frontiers that obtain when we set expected returns equal to their long-term sample means, but we let the variance-covariance of returns change across investment horizons according to the VAR estimates reported in Table 2 and illustrated in Figures 1 and 2. The appendix shows how we perform these computations.

Traditional mean-variance analysis (Markowitz 1952) focuses on risk at short horizons between a month and a year. When the term structure of risk is flat, the efficient frontier is the same at all horizons. Thus short-term mean-variance analysis provides answers that are valid for all mean-variance investors, regardless of their investment horizon. However, when expected returns are time-varying and the term structure of risk is not flat, efficient frontiers at different horizons do not coincide. In that case short-term mean-variance analysis can be misleading for investors with longer investment horizons.

Mean-variance analysis shows that any efficient portfolio is a combination of any two other efficient portfolios. It is standard practice to choose the “global minimum variance portfolio” (GMV portfolio henceforth) as one of those portfolios. This portfolio has intuitive appeal, since it is the portfolio with the smallest variance or risk in the efficient set—the leftmost point in the mean-variance diagram. When a riskless asset (an asset with zero return variance) is available, this portfolio is obviously 100% invested in that asset. When there is no riskless asset, this portfolio is invested in the combination of assets that minimize portfolio return variance regardless of expected return.

Figure 3 plots the annualized standard deviation of the real return on the GMV portfolio implied by the VAR estimates shown in Table 2. For comparison, it also plots the annualized standard deviation of the real return on T-bills, since it is also

standard practice of mean-variance analysis to consider T-bills as a riskless asset, and to take their return as the riskfree rate. Figure 3 shows two results. First, the global minimum variance portfolio is risky at all horizons; second, its risk is similar to the risk of T-bills at short horizons, but is considerably smaller at long horizons.

Figure 4 plots the composition of the GMV portfolio at horizons of 1 quarter, and 5, 10, 25, 50 and 100 years. The results in Figure 3 imply that at long horizons the composition of the GMV portfolio must be different from a 100% T-bill portfolio. Figure 4 shows that the GMV portfolio is fully invested in T-bills at horizons of up to 5 years, but the allocation to bills declines dramatically at longer horizons while the weight of the fixed-maturity 5 year bond increases.⁹ Interestingly, stocks also have a sizable weight in the GMV portfolio at intermediate and long horizons. At horizons of 25 years, long-term bonds already represent about 20% of the GMV portfolio, and stocks represent 12%; at the longest end of the term structure, long-term bonds represent about 62% of the portfolio, and stocks represent 18%, with T-bills completing the remaining 20%.

These results suggest that the standard practice of considering T-bills as the riskless asset works well at short horizons, but that it can be deceptive at long horizons. At short horizons matching their maturity, T-bills carry only short-term inflation risk, which is modest; however, at long horizons they are subject to reinvestment (or real interest) risk, which is important. By contrast, mean-reversion in stock and bond returns makes their volatility decrease with the investment horizon.

If T-bills are not truly riskless even at short horizons, what is the riskless asset? For a short-horizon, buy-and-hold investor, the riskless asset would be a T-bill not subject to inflation risk—an inflation-indexed T-bill. This asset provides a sure real payment at the end of the investor’s horizon. By extension, the riskless asset for a long-term, buy-and-hold investor must then be a zero-coupon inflation-indexed bond whose maturity matches her horizon, since this type of bond provides the investor with a sure cash inflow exactly at the moment the investor needs it. Our empirical analysis suggest that, in the absence of inflation-indexed bonds, the best empirical proxy for this type of bond is a portfolio primarily invested in long-term nominal bonds, plus some stocks and T-bills.

⁹Figure 4 assumes that short positions are allowable, and the GMV portfolio has small short positions in 5-year bonds and stocks for short investment horizons. If we ruled out short positions, the short-horizon GMV portfolio would be fully invested in Treasury bills.

To fully characterize the efficient frontiers at all horizons, we need a second mean-variance efficient portfolio. For comparability with traditional mean-variance analysis, we choose as our second portfolio the tangency portfolio of stocks and the 5-year bond we would obtain if T-bills were truly riskless at all horizons, with a real rate equal to the long-term average shown in Table 1. This would also be the tangency portfolio if inflation-indexed bonds were available and offered this constant real yield. Table 3 shows the composition and Sharpe ratio of the tangency portfolio at horizons of 1, 10, 25 and 100 years.

Table 3 shows that the Sharpe ratio of the tangency portfolio is about 50% larger at a 10-year horizon than at a one-year horizon, and twice as large at a 25-year horizon. At a one-year horizon, the portfolio is invested 47% in equities, and 53% in nominal bonds. The allocation to stocks then increases rapidly, and reaches a maximum weight of about 145% at intermediate horizons of up to 10 years. This is the result of the rapidly declining variance of stock returns (shown in Figure 1), and the rapidly increasing positive correlation between stocks and bonds (shown in Figure 2) that takes place at those horizons. The increasing correlation pushes the portfolio towards the asset with the largest Sharpe ratio, and the declining variance makes this asset even more attractive at those horizons.

At horizons beyond 10 years the allocation to stocks declines, but it stays well above the short-horizon allocation. At the extreme long end of the term structure, the tangency portfolio is still invested 67% in stocks, and 33% in bonds. Once again, Figures 1 and 2 are useful to understand this result. Figure 2 shows that the correlation between stocks and bonds quickly reverts back to levels similar to the short-term correlation for horizons beyond 10 years, but Figure 1 shows that the variance per period of stock returns experiences further reductions. Thus at horizons beyond 10 years, mean-reversion in stock returns is solely responsible for the larger allocation to stocks.

6 Conclusion

This paper has explored the implications for long-term investors of the empirical evidence on the predictability of asset returns. Using a parsimonious yet powerful model of return dynamics, it shows that return forecasting variables such as dividend yields, interest rates, and yield spreads, have substantial effects on optimal portfolio

allocations among bills, stocks, and nominal and inflation-indexed bonds.

For long-horizon, buy-and-hold investors, these effects work through the effect of asset return predictability on the volatility and correlation structure of asset returns across investment horizons, i.e., through the term structure of the risk-return tradeoff. Using data from the U.S. stock, bond, and T-bill markets in the postwar period, the paper fully characterizes the term structure of risk, and shows that the variance and correlation structure of real returns on these assets changes dramatically across investment horizons. These effects reflect underlying changes in stock market risk, inflation risk and real interest risk across investment horizons.

The paper finds that mean-reversion in stock returns decreases the volatility per period of real stock returns at long horizons, while reinvestment risk increases the volatility per period of real T-bill returns. Inflation risk increases the volatility per period of the real return on long-term nominal bonds held to maturity. The paper also finds that stocks and bonds exhibit relatively low positive correlation at both ends of the term structure of risk, but they are highly positively correlated at intermediate investment horizons. Inflation is negatively correlated with bond and stock real returns at short horizons, but positively correlated at long horizons.

These patterns have important implications for the efficient mean-variance frontiers that investors face at different horizons, and suggest that asset allocation recommendations based on short-term risk and return may not be adequate for long horizon investors. For example, the composition of the global minimum variance (GMV) portfolio changes dramatically across investment horizons. We calculate the GMV portfolio when predictor variables are at their unconditional means, that is when market conditions are average, and find that at short horizons it consists almost exclusively of T-bills, but at long horizons reinvestment risk makes T-bills risky, and long-term investors can achieve lower risk with a portfolio that consists predominantly of long-term bonds and stocks.

The paper also finds that the tangency portfolio of bonds and stocks (calculated under the counterfactual assumption that a riskless long-term asset exists with a return equal to the average T-bill return) has a composition that is increasingly biased toward stocks as the horizon increases. This is the result of the increasing positive correlation between stocks and bonds at intermediate horizons, and the decrease of the volatility per period of stock returns at long investment horizons.

It is important to understand that our results depend on the particular model

of asset returns that we have estimated. We have treated the parameters of our VAR(1) model as known, and have studied their implications for long-term portfolio choice. In fact these parameters are highly uncertain, and investors should take this uncertainty into account in their portfolio decisions. A formal Bayesian approach to parameter uncertainty is possible although technically challenging (Xia 2001), but in practice it may be more appealing to study the robustness of portfolio weights to plausible variations in parameters and model specifications. Fortunately the main conclusions discussed here appear to hold up well when the model is estimated over subsamples, or is extended to allow higher-order lags.

The concept of a term structure of the risk-return tradeoff is conceptually appealing but, strictly speaking, is only valid for buy-and-hold investors who make a one-time asset allocation decision and are interested only in the assets available for spending at the end of a particular horizon (Barberis 2000). In practice, however, few investors can truly be characterized as buy-and-hold. Most investors, both individuals and institutions such as pension funds and endowments, can rebalance their portfolios, and have recurrent spending needs which they must finance (completely or partially) off their financial portfolios. These investors may want to rebalance their portfolios in response to changes in investment opportunities.

It is tempting to conclude from this argument that only short-term risk is relevant to the investment decisions of long-horizon investors who can rebalance, whether the risk-return tradeoff changes across investment horizons or not. However, Samuelson (1969), Merton (1969, 1971, 1973) and other financial economists have shown that this conclusion is not correct in general. If interest rates and expected asset returns change over time, risk averse, long-term investors should also be interested in protecting (or hedging) their long-term spending programs against an unexpected deterioration in investment opportunities. Brennan, Schwartz and Lagnado (1997) have coined the term “Strategic Asset Allocation” (SAA) to designate optimal asset allocation rebalancing strategies in the face of changing investment opportunities. SAA portfolios are a combination of two portfolios. The first portfolio is a short-term, mean-variance efficient portfolio. It reflects short-term, or myopic, considerations. The second portfolio, which Merton (1969, 1971, 1973) called the “intertemporal hedging portfolio,” reflects long-term, dynamic hedging considerations.

Using an empirical model for investment opportunities similar to our VAR(1) model, Campbell, Chan and Viceira (2003), Campbell and Viceira (2002) and others have found that asset return predictability can have large effects on the asset alloca-

tion decisions of rebalancing investors. Strategic or intertemporal hedging portfolios tilt the total portfolio away from the short-term mean-variance frontier, as the investor sacrifices some expected portfolio return in exchange for protection from, say, a sudden decrease in expected stock returns or real interest rates.

In contrast to the appealing simplicity of buy-and-hold, mean-variance portfolios, strategic portfolios are in practice difficult to compute, especially as the number of assets and state variables increase. Campbell, Chan and Viceira (2003) and Brandt, Goyal, Santa-Clara and Stroud (2004) have proposed approximate solution methods to compute SAA portfolios. The SAA portfolios calculated by Campbell, Chan, and Viceira have similar qualitative properties to the long-horizon mean-variance portfolios discussed in this paper, but it is important to compare the two approaches more systematically and this is one subject of our ongoing research.

7 References

- Aït-Sahalia, Y., Brandt, M., 2001. Variable selection for portfolio choice. *Journal of Finance* 56, 1297–1351.
- Avramov, D., 2002, Stock return predictability and model uncertainty, *Journal of Financial Economics* 64, 423–458.
- Barberis, N. C., 2000. Investing for the long run when returns are predictable. *Journal of Finance* 55, 225–264.
- Bekaert, G., Hodrick, R. J., Marshall, D. A., 1997. On biases in tests of the expectations hypothesis of the term structure of interest rates. *Journal of Financial Economics* 44, 309–348.
- Bollerslev, Tim, 1990. Modeling the coherence in short run nominal exchange rates: A multivariate generalized ARCH model. *Review of Economics and Statistics* 72, 498–505.
- Brandt, M., Goyal, A., Santa-Clara, P., and Stroud, J.R., 2004. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies*, forthcoming.
- Brennan, M. J., Schwartz, E. S., Lagnado, R., 1997. Strategic asset allocation. *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Campbell, J. Y., 1987. Stock returns and the term structure. *Journal of Financial Economics* 18, 373–399.
- Campbell, J. Y., 1991. A variance decomposition for stock returns. *Economic Journal* 101, 157–179.
- Campbell, J.Y., 2001. Why long horizons? A study of power against persistent alternatives. *Journal of Empirical Finance* 8, 459–491.
- 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., 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1, 195–228.

- Campbell, J. Y., Shiller, R. J., 1991. Yield spreads and interest rates: a bird's eye view. *Review of Economic Studies* 58, 495–514.
- Campbell, J. Y., Viceira, L. M., 1999. Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics* 114, 433–495.
- Campbell, J. Y., Viceira, L. M., 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press, Oxford.
- Campbell, J.Y., Vuolteenaho, T., 2004. Inflation illusion and stock prices. *American Economic Review, Papers and Proceedings*, 94, 19–23.
- Campbell, J. Y., Yogo, M., 2004. Efficient tests of stock return predictability. Manuscript, Harvard University.
- Cavanagh, C., Elliott, G., Stock, J.H., 1995. Inference in models with nearly integrated regressors. *Econometric Theory*, 1131-1147.
- Chacko, G., Viceira, L. M., 1999. Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. NBER Working Paper 7377. National Bureau of Economic Research, Cambridge, MA.
- Elliott, G., and Stock, J.H., 1994. Inference in times series regression when the order of integration of a regressor is unknown. *Econometric Theory* 10, 672-700.
- Engle, Robert, 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics* 20, 339-350.
- Fama, E., 1984. The information in the term structure. *Journal of Financial Economics* 13, 509–528.
- Fama, E., French, K., 1988. Dividend yields and expected stock returns. *Journal of Financial Economics* 22, 3–27.
- Fama, E., French, K., 1989. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25, 23–49.
- Fama, E., Schwert, G. W., 1977. Asset returns and inflation. *Journal of Financial Economics* 5, 115–146.

- Glosten, L. R., Jagannathan, R., Runkle, D., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779–1801.
- Goetzmann, W. N., Jorion, P., 1993. Testing the predictive power of dividend yields. *Journal of Finance* 48, 663–679.
- Harvey, C., 1989. Time-varying conditional covariances in tests of asset pricing models. *Journal of Financial Economics* 22, 305–334.
- Harvey, C., 1991. The world price of covariance risk. *Journal of Finance* 46, 111–157.
- Hodrick, R. J., 1992. Dividend yields and expected stock returns: alternative procedures for inference and measurement. *Review of Financial Studies* 5, 357–386.
- Hoevenaars R.P.M.M., Molenaar R.D.J, and Steenkamp T.B.M., 2004. Strategic asset allocation with parameter uncertainty. Working paper, Maastricht University, ABP, Vrije University.
- Kandel, S., Stambaugh, R., 1987. Long horizon returns and short horizon models. CRSP Working Paper No.222. University of Chicago.
- Markowitz, H., 1952. Portfolio selection. *Journal of Finance* 7, 77–91.
- 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, 867–87.
- Modigliani, F., Cohn, R., 1979. Inflation, rational valuation, and the market. *Financial Analysts Journal*.
- Nelson, C. R., Kim, M. J., 1993. Predictable stock returns: the role of small sample bias. *Journal of Finance* 48, 641–661.
- Rigobon, R., Sack, B., 2003. Spillovers across U.S. financial markets. NBER Working Paper 9640. National Bureau of Economic Research, Cambridge, MA.

- Ritter, J.R., and Warr, R.S., 2002. The decline of inflation and the bull market of 1982-1999. *Journal of Financial and Quantitative Analysis* 37, 29-61.
- Samuelson, P. A., 1969. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics* 51, 239-246.
- Shiller, R. J., Campbell, J. Y., Schoenholtz, K. L., 1983. Forward rates and future policy: interpreting the term structure of interest rates. *Brookings Papers on Economic Activity* 1, 173-217.
- Siegel, Jeremy, 1994. *Stocks for the Long Run*, McGraw-Hill, New York, NY.
- Stambaugh, R. F., 1999. Predictive regressions. *Journal of Financial Economics* 54, 375-421.
- Xia, Y., 2001. Learning about predictability: The effect of parameter uncertainty on dynamic asset allocation. *Journal of Finance* 56, 205-246.

8 Appendix. Long-Horizon Mean-Variance Analysis: A User Guide

8.1 Long-horizon moments in a VAR(1)

This section derives expressions for the conditional mean and variance-covariance matrix of $(\mathbf{z}_{t+1} + \dots + \mathbf{z}_{t+k})$ for any horizon k , where the vector \mathbf{z}_{t+1} follows the VAR(1) process described in Section 2.¹⁰ This is important for our subsequent portfolio analysis across investment horizons, because returns are a subset of the vector \mathbf{z}_{t+1} , and cumulative k -period log returns are obtained by adding one-period log returns over k successive periods.

8.1.1 Conditional k-period moments

We start by deriving a set of equations that relate future values of the vector of state variables to its current value \mathbf{z}_t plus a weighted sum of interim shocks:

$$\begin{aligned}
 z_{t+1} &= \Phi_0 + \Phi_1 z_t + v_{t+1} \\
 z_{t+2} &= \Phi_0 + \Phi_1 z_{t+1} + v_{t+2} \\
 &= \Phi_0 + \Phi_1 \Phi_0 + \Phi_1 \Phi_1 z_t + \Phi_1 v_{t+1} + v_{t+2} \\
 z_{t+3} &= \Phi_0 + \Phi_1 z_{t+2} + v_{t+3} \\
 &= \Phi_0 + \Phi_1 \Phi_0 + \Phi_1 \Phi_1 \Phi_0 + \Phi_1 \Phi_1 \Phi_1 z_t + \Phi_1 \Phi_1 v_{t+1} + \Phi_1 v_{t+2} + v_{t+3} \\
 z_{t+k} &= \Phi_0 + \Phi_1 \Phi_0 + \Phi_1^2 \Phi_0 + \dots + \Phi_1^{k-1} \Phi_0 + \Phi_1^k z_t \\
 &\quad + \Phi_1^{k-1} v_{t+1} + \Phi_1^{k-2} v_{t+2} + \dots + \Phi_1 v_{t+k-1} + v_{t+k}
 \end{aligned}$$

These expressions follow immediately from forward recursion of equation (3).

¹⁰Avramov (2002) also derives formulas for the conditional mean and variance of returns over long horizons in a framework similar to ours.

Adding the expressions for z_{t+1}, z_{t+1}, \dots , and reordering terms yields:

$$\begin{aligned}
z_{t+1} + \dots + z_{t+k} &= [k + (k-1)\Phi_1 + (k-2)\Phi_1^2 + \dots + \Phi_1^{k-1}]\Phi_0 \\
&\quad + (\Phi_1^k + \Phi_1^{k-1} + \dots + \Phi_1)z_t \\
&\quad + (1 + \Phi_1 + \dots + \Phi_1^{k-1})v_{t+1} \\
&\quad + (1 + \Phi_1 + \dots + \Phi_1^{k-2})v_{t+2} \\
&\quad + \dots \\
&\quad + (I + \Phi_1)v_{t+k-1} + v_{t+k}.
\end{aligned}$$

We can write this expression more compactly as

$$z_{t+1} + \dots + z_{t+k} = \left[\sum_{i=0}^{k-1} (k-i)\Phi_1^i \right] \Phi_0 + \left[\sum_{j=1}^k \Phi_1^j \right] z_t + \sum_{q=1}^k \left[\sum_{p=0}^{k-q} \Phi_1^p v_{t+q} \right].$$

Now we are ready to compute conditional k-period moments of the state vector. The conditional mean is given by

$$\mathbb{E}_t(z_{t+1} + \dots + z_{t+k}) = \left[\sum_{i=0}^{k-1} (k-i)\Phi_1^i \right] \Phi_0 + \left[\sum_{j=1}^k \Phi_1^j \right] z_t,$$

since the shocks v_{t+q} have zero mean.

The conditional variance is given by

$$\begin{aligned}
\text{Var}_t(z_{t+1} + \dots + z_{t+k}) &= \text{Var}_t \left[\left[\sum_{i=0}^{k-1} (k-i)\Phi_1^i \right] \Phi_0 + \left[\sum_{j=1}^k \Phi_1^j \right] z_t + \sum_{q=1}^k \left[\sum_{p=0}^{k-q} \Phi_1^p v_{t+q} \right] \right] \\
&= \text{Var}_t \left[\sum_{q=1}^k \left[\sum_{p=0}^{k-q} \Phi_1^p v_{t+q} \right] \right],
\end{aligned}$$

since all other terms are constant or known at time t. Expanding this expression we

obtain:

$$\begin{aligned}
\text{Var}_t(z_{t+1} + \dots + z_{t+k}) &= \text{Var}_t[(I + \Phi_1 + \dots + \Phi_1^{k-1})v_{t+1} + (I + \Phi_1 + \dots + \Phi_1^{k-2})v_{t+2} \\
&\quad + \dots + (I + \Phi_1)v_{t+k-1} + v_{t+k}] \\
&= \Sigma_v + (I + \Phi_1)\Sigma_v(I + \Phi_1)' \\
&\quad + (I + \Phi_1 + \Phi_1\Phi_1)\Sigma_v(I + \Phi_1 + \Phi_1\Phi_1)' \\
&\quad + \dots \\
&\quad + (I + \Phi_1 + \dots + \Phi_1^{k-1})\Sigma_v(I + \Phi_1' + \dots + \Phi_1^{k-1})',
\end{aligned}$$

which follows from reordering terms and noting that the conditional variance-covariance matrix of v_{t+i} is the same (Σ_v) at all leads i .

8.1.2 Conditional k -period moments of returns

We are only interested in extracting conditional moments per period from the portion of the VAR that contains returns. We can do so by using selector matrices. For example, if we want to recover the annualized (or per period) k -period conditional moments of the benchmark asset return and excess returns, we can use the following matrix:

$$\mathbf{H}_r = \begin{bmatrix} \mathbf{I}_{(n+1) \times (n+1)} & \mathbf{0}_{(m-n) \times (m-n)} \end{bmatrix}.$$

It is straightforward to check that this matrix selects the k -period expected return and variance-covariance matrix of returns when applied to the expressions for the conditional expectation and variance-covariance matrix of $(\mathbf{z}_{t+1} + \dots + \mathbf{z}_{t+k})$.

Thus we can extract the vector of expected k -period returns per period as follows:

$$\frac{1}{k} \begin{bmatrix} \text{E}_t \left(r_{0,t+1}^{(k)} \right) \\ \text{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \right) \end{bmatrix} = \frac{1}{k} \mathbf{H}_r \text{E}_t(z_{t+1} + \dots + z_{t+k}),$$

and the k -period return variance-covariance matrix per period as

$$\frac{1}{k} \text{Var}_t \begin{bmatrix} r_{0,t+1}^{(k)} \\ \mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \end{bmatrix} = \frac{1}{k} \begin{bmatrix} \sigma_0^2(k) & \boldsymbol{\sigma}_{0x}(k)' \\ \boldsymbol{\sigma}_{0x}(k) & \boldsymbol{\Sigma}_{xx}(k) \end{bmatrix} = \frac{1}{k} \mathbf{H}_r \text{Var}_t(z_{t+1} + \dots + z_{t+k}) \mathbf{H}_r'.$$

Similarly, the following selector matrix extracts per-period, k -period conditional moments of log real returns:

$$\mathbf{M}_r = \begin{bmatrix} 1 & \mathbf{0}_{1 \times n} & \mathbf{0}_{1 \times (m-n-1)} \\ \boldsymbol{\iota}_{n \times n} & \mathbf{I}_{n \times n} & \mathbf{0}_{(m-n-1) \times (m-n-1)} \end{bmatrix}, \quad (17)$$

which implies

$$\frac{1}{k} \begin{bmatrix} \mathbb{E}_t \left(r_{0,t+1}^{(k)} \right) \\ \mathbb{E}_t \left(\mathbf{r}_{t+1}^{(k)} \right) \end{bmatrix} = \frac{1}{k} \mathbf{M}_r \mathbb{E}_t (z_{t+1} + \dots + z_{t+k})$$

and

$$\frac{1}{k} \text{Var}_t \begin{bmatrix} r_{0,t+1}^{(k)} \\ \mathbf{r}_{t+1}^{(k)} \end{bmatrix} = \frac{1}{k} \boldsymbol{\Sigma}_{rr} (k) = \frac{1}{k} \mathbf{M}_r \text{Var}_t (z_{t+1} + \dots + z_{t+k}) \mathbf{M}_r'.$$

8.2 One-Period Mean-Variance Analysis

This section derives the standard mean-variance for one-period log returns. The next section extends these results to a multi-period environment. We start by noting that the gross return on the wealth portfolio is given by

$$\begin{aligned} R_{p,t+1} &= \sum_{i=1}^n \alpha_{i,t} (R_{i,t+1} - R_{0,t+1}) + R_{0,t+1} \\ &= \boldsymbol{\alpha}'_t (\mathbf{R}_{t+1} - R_{0,t+1} \boldsymbol{\iota}) + R_{0,t+1}, \end{aligned} \quad (18)$$

where R_i is the gross return on asset i , R_0 is the gross return on the benchmark asset, and $\alpha_{i,t}$ is the portfolio weight on asset i ; $\boldsymbol{\alpha}_t$ is a $(n \times 1)$ vector containing the portfolio weights $\alpha_{i,t}$'s, and $\boldsymbol{\iota}$ is a $(n \times 1)$ vector of 1's.

The return on the portfolio in equation (18) is expressed in terms of the gross returns on the assets. Since we are interested in working with log returns, we need to derive an expression for the log return on the portfolio. Campbell and Viceira (1999, 2001) and Campbell, Chan and Viceira (2003) suggest the following approximation to the log return on the portfolio:

$$r_{p,t+1} = r_{0,t+1} + \boldsymbol{\alpha}'_t (\mathbf{r}_{t+1} - r_{0,t+1} \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\alpha}'_t (\boldsymbol{\sigma}_x^2 - \boldsymbol{\Sigma}_{xx} \boldsymbol{\alpha}_t), \quad (19)$$

where

$$\boldsymbol{\Sigma}_{xx} \equiv \text{Var}_t (\mathbf{r}_{t+1} - r_{0,t+1} \boldsymbol{\iota}),$$

and

$$\boldsymbol{\sigma}_x^2 \equiv \text{diag}(\boldsymbol{\Sigma}_{xx}).$$

We now derive the one-period mean-variance frontier with and without a riskless asset, and the weights for the “tangency portfolio,” i.e. the portfolio containing only risky assets that also belongs simultaneously to both the mean-variance frontier of risky assets, and the mean-variance frontier with a riskless asset.

8.2.1 One-Period Mean-Variance Frontier with Log Returns and No Riskless Asset

This section derives the mean-variance frontier when none of the assets available for investment is unconditionally riskless in real terms at a one-period horizon. This happens, for example, when there is significant short-term inflation risk, and the only short-term assets available for investment are nominal money market instruments, such as nominal T-bills. In practice short-term inflation risk is small, and one can probably ignore it. However, we still think it is worth analyzing the case with no one-period riskless asset, to gain intuition for the subsequent results for long horizons, where inflation risk is significant, and nominal bonds can be highly risky in real terms.

Problem We state the mean-variance portfolio optimization problem with log returns as follows:

$$\min_{\boldsymbol{\alpha}} \frac{1}{2} \text{Var}_t(r_{p,t+1}), \quad (20)$$

subject to

$$\text{E}_t(r_{p,t+1}) + \frac{1}{2} \text{Var}_t(r_{p,t+1}) = \mu_p,$$

where $r_{p,t+1}$ is given in equation (19). That is, we want to find the vector of portfolio weights $\boldsymbol{\alpha}_t$ that minimizes the variance of the portfolio log return when the required expected gross return on the portfolio is μ_p .

Equation (19) implies that the variance of the log portfolio return is given by

$$\text{Var}_t(r_{p,t+1}) = \boldsymbol{\alpha}'_t \boldsymbol{\Sigma}_{xx} \boldsymbol{\alpha}_t + \sigma_0^2 + 2\boldsymbol{\alpha}'_t \boldsymbol{\sigma}_{0x}$$

and the log expected portfolio return is given by

$$\mathbb{E}_t(r_{p,t+1}) + \frac{1}{2} \text{Var}_t(r_{p,t+1}) = \left(\mathbb{E}_t r_{0,t+1} + \frac{1}{2} \sigma_0^2 \right) + \boldsymbol{\alpha}'_t \left(\mathbb{E}_t(\mathbf{r}_{t+1} - r_{0,t+1} \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_x^2 + \boldsymbol{\sigma}_{0x} \right). \quad (21)$$

It is useful to note here that the VAR(1) model for returns given in equation (3) implies that

$$\mathbb{E}_t r_{0,t+1} = \mathbf{H}'_0 (\boldsymbol{\Phi}_0 + \boldsymbol{\Phi}_1 \mathbf{z}_t)$$

and

$$\begin{aligned} \mathbb{E}_t(\mathbf{r}_{t+1} - r_{0,t+1} \boldsymbol{\iota}) &= \mathbb{E}_t(\mathbf{x}_{t+1}) \\ &= \mathbf{H}_x (\boldsymbol{\Phi}_0 + \boldsymbol{\Phi}_1 \mathbf{z}_t), \end{aligned}$$

where \mathbf{H}'_0 is a $(1 \times m)$ selection vector $(1, 0, \dots, 0)$ that selects the first element of the matrix which it multiplies, and \mathbf{H}_x is a $(n \times m)$ selection matrix that selects rows 2 through $n + 1$ of the matrix which it multiplies:

$$\mathbf{H}_x = \begin{bmatrix} \mathbf{0}_{n \times 1} & \mathbf{I}_{n \times n} & \mathbf{0}_{(m-n-1) \times (m-n-1)} \end{bmatrix},$$

where $\mathbf{I}_{n \times n}$ is an identity matrix.

The Lagrangian for this problem is

$$\mathcal{L} = \frac{1}{2} \text{Var}_t(r_{p,t+1}) + \lambda \left(\mu_P - \left(\mathbb{E}_t(r_{p,t+1}) + \frac{1}{2} \text{Var}_t(r_{p,t+1}) \right) \right),$$

where λ is the Lagrange multiplier on the expected portfolio return constraint. Note that our definition of the portfolio return already imposes the constraint that portfolio weights must add up to one.

Mean-variance portfolio rule The solution to the mean-variance program (20) is

$$\boldsymbol{\alpha}_t = \lambda \boldsymbol{\Sigma}_{xx}^{-1} \left[\mathbb{E}_t(\mathbf{r}_{t+1} - r_{0,t+1} \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_x^2 + \boldsymbol{\sigma}_{0x} \right] - \boldsymbol{\Sigma}_{xx}^{-1} \boldsymbol{\sigma}_{0x}. \quad (22)$$

We can rewrite $\boldsymbol{\alpha}_t$ as a portfolio combining two distinct portfolios:

$$\boldsymbol{\alpha}_t = \lambda \boldsymbol{\Sigma}_{xx}^{-1} \left[\mathbb{E}_t(\mathbf{r}_{t+1} - r_{0,t+1} \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_x^2 \right] + (1 - \lambda) (-\boldsymbol{\Sigma}_{xx}^{-1} \boldsymbol{\sigma}_{0x}).$$

The first portfolio is proportional to the expected gross excess return on all assets, weighted by the inverse of the variance-covariance matrix of excess returns. If asset returns are independent of each other, so Σ_{xx} is a diagonal matrix containing the variances of asset returns, this portfolio is simply the vector of Sharpe ratios for each asset class—each one divided by its return standard deviation.

The second portfolio, with weights $\Sigma_{xx}^{-1}\sigma_{0x}$, is the portfolio of assets with minimum absolute variance, or global minimum-variance portfolio. That is, this portfolio is the solution to:

$$-\Sigma_{xx}^{-1}\sigma_{0x} = \min_{\alpha} \frac{1}{2} \text{Var}_t(r_{p,t+1}).$$

In the literature on dynamic optimal portfolio choice, the mean-variance portfolio is also known as the “myopic portfolio.” This name derives from the fact that it takes into consideration only expected returns next period, ignoring other considerations that might affect the portfolio decisions of risk averse investors, such as the ability that certain asset classes might have to protect the spending needs of these investors from a sudden deterioration in investment opportunities.

Equation for mean-variance frontier Equation (22) implies the following relation between $\sigma_p^2 \equiv \text{Var}_t(r_{p,t+1})$ and μ_p .¹¹

$$\sigma_p^2 = \frac{1}{A} (e\mu_p^2 + 2Be\mu_p + B^2) + \sigma_0^2 - C, \quad (23)$$

where $e\mu_p$ is the expected excess return on the portfolio over the benchmark asset,

$$e\mu_p = \mu_p - \left(\text{E}_t r_{0,t+1} + \frac{1}{2}\sigma_0^2 \right),$$

and

$$A = \left(\text{E}_t (\mathbf{r}_{t+1} - r_{0,t+1}\mathbf{l}) + \frac{1}{2}\sigma_x^2 + \sigma_{0x} \right)' \Sigma_{xx}^{-1} \left(\text{E}_t (\mathbf{r}_{t+1} - r_{0,t+1}\mathbf{l}) + \frac{1}{2}\sigma_x^2 + \sigma_{0x} \right),$$

$$B = \sigma'_{0x} \Sigma_{xx}^{-1} \left(\text{E}_t (\mathbf{r}_{t+1} - r_{0,t+1}\mathbf{l}) + \frac{1}{2}\sigma_x^2 + \sigma_{0x} \right),$$

$$C = \sigma'_{0x} \Sigma_{xx}^{-1} \sigma_{0x}.$$

¹¹To get this equation, note that at the optimum either $\lambda = 0$, or $\mu_p = (\text{E}_t(r_{p,t+1}) + \frac{1}{2}\sigma_p^2)$.

8.2.2 One-Period Mean-Variance Frontier with Log Returns and a Riskless Asset

In this section we consider mean-variance analysis with log returns when investors have an additional asset which is truly riskless in real terms over one period. This asset is an inflation-indexed T-Bill. We denote its return as r_f . Thus the investor has now $n + 2$ assets available for investment. This derivation is useful to compute the tangency portfolio.

Problem

$$\begin{aligned} \min_{\alpha} \frac{1}{2} \text{Var}_t (r_{p,t+1} - r_f) \\ \text{subject to } \text{E}_t (r_{p,t+1} - r_f) + \frac{1}{2} \text{Var}_t (r_{p,t+1} - r_f) = \mu_p - r_f, \end{aligned} \quad (24)$$

where

$$r_{p,t+1} - r_f = \boldsymbol{\omega}'_t (\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\omega}'_t (\boldsymbol{\sigma}_r^2 - \boldsymbol{\Sigma}_{rr} \boldsymbol{\omega}_t),$$

$\boldsymbol{\omega}_t$ is a $(n + 1) \times 1$ vector of portfolio weights,

$$\boldsymbol{\Sigma}_{rr} \equiv \text{Var}_t (\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}),$$

and

$$\boldsymbol{\sigma}_r^2 \equiv \text{diag} (\boldsymbol{\Sigma}_{rr}).$$

The variance of the log excess portfolio return is given by

$$\text{Var}_t (r_{p,t+1} - r_f) = \boldsymbol{\omega}'_t \boldsymbol{\Sigma}_{rr} \boldsymbol{\omega}_t \quad (25)$$

and the log excess expected portfolio return is given by:

$$\text{E}_t (r_{p,t+1} - r_f) + \frac{1}{2} \text{Var}_t (r_{p,t+1} - r_f) = \boldsymbol{\omega}'_t \left(\text{E}_t (\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_r^2 \right). \quad (26)$$

Note that:

$$\text{E}_t (\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) = \mathbf{M}_r (\boldsymbol{\Phi}_0 + \boldsymbol{\Phi}_1 \mathbf{z}_t) - r_f \boldsymbol{\iota},$$

where \mathbf{M}_r is the $(n + 1 \times m)$ selection matrix defined in (17).

The Lagrangian for this problem is

$$\mathcal{L} = \frac{1}{2} \text{Var}_t(r_{p,t+1} - r_f) + \lambda_f \left(\mu_P - r_f - \left(\mathbf{E}_t(r_{p,t+1} - r_f) + \frac{1}{2} \text{Var}_t(r_{p,t+1} - r_f) \right) \right),$$

where λ is the Lagrange multiplier on the expected portfolio return constraint.

Mean-variance portfolio rule The solution to the mean-variance program (24) is

$$\boldsymbol{\omega}_t = \lambda_f \boldsymbol{\Sigma}_{rr}^{-1} \left[\mathbf{E}_t(\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_x^2 \right]. \quad (27)$$

Equation for mean-variance frontier Equation (27) implies the following relation between $\sigma_p^2 \equiv (\text{Var}_t r_{p,t+1} - r_f)$ and μ_p :

$$\sigma_p^2 = \frac{1}{A} e\mu_p^2, \quad (28)$$

where $e\mu_p$ is the expected excess return on the portfolio over the benchmark asset,

$$e\mu_p = \mu_p - r_f$$

and

$$A = \left(\mathbf{E}_t(\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_r^2 \right)' \boldsymbol{\Sigma}_{xx}^{-1} \left(\mathbf{E}_t(\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_r^2 \right).$$

Equation (28) implies that, in a standard deviation-expected return space, the frontier is a line that intersects the expected return axis at r_f , and it has a slope equal to \sqrt{A} . We show below that this slope is the Sharpe ratio of the tangency portfolio.

Tangency portfolio We can now derive the weights for the tangency portfolio between frontiers (23) and (28). This is a portfolio whose weights $\boldsymbol{\omega}_{T,t}$ satisfy $\boldsymbol{\omega}'_{T,t} \boldsymbol{\iota} = \mathbf{1}$. That is, it is a portfolio with no loading on the risk-free asset.

We can easily solve for the weights of this portfolio by finding the Lagrange multiplier λ_f for which $\boldsymbol{\omega}_t$ in (27) satisfies the constraint $\boldsymbol{\omega}'_{T,t} \boldsymbol{\iota} = \mathbf{1}$:

$$\mathbf{1} = \boldsymbol{\omega}'_{T,t} \boldsymbol{\iota} = \lambda_f \left[\mathbf{E}_t(\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_r^2 \right]' (\boldsymbol{\Sigma}_{rr}^{-1})' \boldsymbol{\iota},$$

which implies

$$\lambda_f = \frac{1}{[\mathbf{E}_t(\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_r^2]' \boldsymbol{\Sigma}_{rr}^{-1} \boldsymbol{\iota}}.$$

Therefore,

$$\boldsymbol{\omega}_{T,t} = \frac{1}{[\mathbf{E}_t(\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_r^2]' (\boldsymbol{\Sigma}_{rr}^{-1})' \boldsymbol{\iota}} \boldsymbol{\Sigma}_{rr}^{-1} \left[\mathbf{E}_t(\mathbf{r}_{t+1} - r_f \boldsymbol{\iota}) + \frac{1}{2} \boldsymbol{\sigma}_r^2 \right]. \quad (29)$$

It is straightforward to show that the Sharpe ratio of the tangency portfolio equals \sqrt{A} , i.e., that the Sharpe ratio is the slope of the mean-standard deviation frontier with a riskless asset. To see this, substitute (29) into equations (25) and (26) to compute expressions for the variance and expected excess return on the tangency portfolio, and compute the ratio of the expected excess return to the standard deviation.

8.3 Long-Horizon Mean-Variance Analysis

In this section we extend to a multiperiod horizon the results of Section 3 for mean-variance efficient allocations and frontiers. We assume that the investor has an investment horizon of k periods, and that she chooses portfolio weights only every k periods, i.e., at $t, t+k, t+2k, \dots$. We denote these weights as $\boldsymbol{\alpha}_t(k)$. The investor chooses those weights that minimize the variance per period of the log portfolio return subject to a required expected return per period on the portfolio.

To solve this problem we need to compute the log k -period portfolio return $r_{p,t+k}^{(k)}$. One possibility is to use, for similarity with the approximation for the log 1-period portfolio return given in (19), the following expression:

$$r_{p,t+k}^{(k)} = r_{0,t+k}^{(k)} + \boldsymbol{\alpha}'_t(k) \left(\mathbf{r}_{t+k}^{(k)} - r_{0,t+k}^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\alpha}'_t(k) \left(\boldsymbol{\sigma}_x^2(k) - \boldsymbol{\Sigma}_{xx}(k) \boldsymbol{\alpha}_t(k) \right), \quad (30)$$

where $\boldsymbol{\iota}$ is a $(n \times 1)$ vector of 1's,

$$\boldsymbol{\Sigma}_{xx}(k) \equiv \text{Var}_t \left(\mathbf{r}_{t+k}^{(k)} - r_{0,t+k}^{(k)} \boldsymbol{\iota} \right),$$

$$\boldsymbol{\sigma}_x^2(k) \equiv \text{diag}(\boldsymbol{\Sigma}_{xx}(k)).$$

and, for any asset (or portfolio of assets) l , the cumulative k -period log return (or log excess return) is

$$r_{l,t+k}^{(k)} = \sum_{i=1}^k r_{l,t+i}. \quad (31)$$

Note that to simplify notation we are ignoring the subscript t when denoting conditional variances and covariances. Our assumption that the dynamics of state variables and returns follow a homoskedastic VAR implies that conditional variances are not a function of time—though they are a function of k , the investment horizon.

It is also important to keep in mind that (30) is an approximation to the log return on wealth. This accuracy of this approximation is high at short horizons, but it deteriorates as the horizon k becomes longer.

8.3.1 K-Period Mean-Variance Frontier with Log Returns and No Riskless Asset

Problem The investor chooses those weights that minimize the variance per period of the log portfolio return subject to a required expected return per period on the portfolio. If no riskless asset over horizon k is available, the problem becomes

$$\begin{aligned} \min_{\alpha} \quad & \frac{1}{2} \frac{\text{Var}_t \left(r_{p,t+k}^{(k)} \right)}{k} \\ \text{subject to} \quad & \frac{\text{E}_t \left(r_{p,t+k}^{(k)} \right) + \frac{1}{2} \text{Var}_t \left(r_{p,t+k}^{(k)} \right)}{k} = \mu_p, \end{aligned} \quad (32)$$

where μ_p is the required portfolio return per period.

>From equation (30), the conditional variance of the cumulative log portfolio return is given by

$$\text{Var}_t \left(r_{p,t+k}^{(k)} \right) = \alpha'_t(k) \Sigma_{xx}(k) \alpha_t(k) + \sigma_0^2(k) + 2\alpha'_t(k) \sigma_{0x}(k),$$

and the cumulative log expected portfolio return is given by

$$\text{E}_t \left(r_{p,t+k}^{(k)} \right) + \frac{1}{2} \sigma_p^2(k) = \left(\text{E}_t r_{0,t+1}^{(k)} + \frac{1}{2} \sigma_0^2(k) \right) + \alpha'_t \left(\text{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \sigma_x^2(k) + \sigma_{0x}(k) \right).$$

Note that we already have expressions for the moments in the RHS of these equations based on our VAR specification.

The Lagrangian for this problem is

$$\mathcal{L} = \frac{1}{2} \frac{\text{Var}_t \left(r_{p,t+k}^{(k)} \right)}{k} + \lambda \left(\mu_P - \frac{\text{E}_t \left(r_{p,t+k}^{(k)} \right) + \frac{1}{2} \text{Var}_t \left(r_{p,t+k}^{(k)} \right)}{k} \right),$$

where λ is the Lagrange multiplier on the expected portfolio return constraint. Note that we have already imposed the constraint that portfolio weights must add up to one in the way we define the log return on the portfolio.

Myopic portfolio rule The solution to the mean-variance program (20) is

$$\boldsymbol{\alpha}_t(k) = \lambda \boldsymbol{\Sigma}_{xx}^{-1}(k) \left[\text{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_x^2(k) + \boldsymbol{\sigma}_{0x}(k) \right] - \boldsymbol{\Sigma}_{xx}^{-1}(k) \boldsymbol{\sigma}_{0x}(k). \quad (33)$$

Again, it is useful to note that we can write (33) as

$$\boldsymbol{\alpha}_t(k) = \lambda \boldsymbol{\Sigma}_{xx}^{-1}(k) \left[\text{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_x^2(k) \right] + (1 - \lambda) \left(-\boldsymbol{\Sigma}_{xx}^{-1}(k) \boldsymbol{\sigma}_{0x}(k) \right),$$

where $\boldsymbol{\Sigma}_{xx}^{-1}(k) \boldsymbol{\sigma}_{0x}(k)$ is the k -horizon global minimum-variance portfolio:

$$-\boldsymbol{\Sigma}_{xx}^{-1}(k) \boldsymbol{\sigma}_{0x}(k) = \min_{\boldsymbol{\alpha}} \frac{1}{2} \text{Var}_t \left(r_{p,t+k}^k \right).$$

Thus $\boldsymbol{\alpha}_t(k)$ is a linear combination of two portfolios, one of which is the global minimum-variance portfolio.

Equation for k-period mean-variance frontier Equation (22) implies the following relation between $\sigma_p^2(k)/k \equiv \text{Var}_t \left(r_{p,t+k}^{(k)} \right)/k$ and μ_p :

$$\frac{\sigma_p^2(k)}{k} = \frac{1}{A(k)} \left(e\mu_p^2 + 2B(k)e\mu_p + B^2(k) \right) + \sigma_0^2(k) - C(k), \quad (34)$$

where $e\mu_p$ is the required expected excess return per period on the portfolio over the benchmark asset,

$$e\mu_p = \mu_p - \frac{1}{k} \left(\mathbb{E}_t r_{0,t+1}^{(k)} + \frac{1}{2} \sigma_0^2(k) \right),$$

and

$$\begin{aligned} A(k) &= \frac{1}{k} \left(\mathbb{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_x^2(k) + \boldsymbol{\sigma}_{0x}(k) \right)' \\ &\quad \times \boldsymbol{\Sigma}_{xx}^{-1}(k) \left(\mathbb{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_x^2(k) + \boldsymbol{\sigma}_{0x}(k) \right), \end{aligned}$$

$$B(k) = \frac{1}{k} \boldsymbol{\sigma}'_{0x}(k) \boldsymbol{\Sigma}_{xx}^{-1}(k) \left(\mathbb{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_{0,t+1}^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_x^2(k) + \boldsymbol{\sigma}_{0x}(k) \right),$$

and

$$C(k) = \frac{1}{k} \boldsymbol{\sigma}'_{0x}(k) \boldsymbol{\Sigma}_{xx}^{-1}(k) \boldsymbol{\sigma}_{0x}(k).$$

Unconditional mean-variance frontier The unconditional mean variance frontier simply follows from taking limits in (34) as $k \rightarrow \infty$. That is, we use the unconditional variance-covariance matrix of log returns implied by the VAR.

8.3.2 K-Period Mean-Variance Frontier with Log Returns and a Riskless Asset

In this section we consider mean-variance analysis with log returns when investors have an additional asset which is truly riskless in real terms over their investment horizon. This asset is an inflation-indexed zero-coupon bond. The investor has now $n+1$ assets available for investment. This derivation is useful to compute the tangency portfolio.

>From the perspective of a k -period mean-variance investor—who is a buy-and-hold investor—the return on the riskfree asset is the yield on this bond, which we denote by $r_f^{(k)}$. We assume that $r_f^{(k)}$ is continuously-compounded, in natural units per period. Thus to compound k periods ahead we need to multiply it by k .

Problem

$$\min_{\alpha} \frac{1}{2} \frac{\text{Var}_t \left(r_{p,t+k}^{(k)} - kr_f^{(k)} \right)}{k} \quad (35)$$

subject to $\frac{\text{E}_t \left(r_{p,t+k}^{(k)} - r_f^{(k)} k \right) + \frac{1}{2} \text{Var}_t \left(r_{p,t+k}^{(k)} - kr_f^{(k)} \right)}{k} = \mu_p - r_f^{(k)},$

where we assume that $r_f^{(k)}$ is given in units per period, and $r_{l,t+k}^{(k)}$ is defined in (31).

To solve this problem we need to compute the log k-period portfolio return $r_{p,t+k}^{(k)}$. We adopt the same convention as in the previous section, and use the k-period ahead conditional covariances for the Jensen's inequality correction in the approximation to the log excess return on wealth k periods ahead:

$$r_{p,t+k}^{(k)} - kr_f^{(k)} = \boldsymbol{\omega}'_t(k) \left(\mathbf{r}_{t+k}^{(k)} - kr_f^{(k)} \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\omega}'_t(k) \left(\boldsymbol{\sigma}_r^2(k) - \boldsymbol{\Sigma}_{rr}(k) \boldsymbol{\omega}_t(k) \right), \quad (36)$$

where

$$\boldsymbol{\Sigma}_{rr}(k) \equiv \text{Var}_t \left(\mathbf{r}_{t+k}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right),$$

$$\boldsymbol{\sigma}_r^2(k) \equiv \text{diag} \left(\boldsymbol{\Sigma}_{rr}(k) \right).$$

Using equation (36), the conditional variance of the cumulative log portfolio return is given by

$$\text{Var}_t \left(r_{p,t+k}^{(k)} - kr_f^{(k)} \right) = \boldsymbol{\omega}'_t(k) \boldsymbol{\Sigma}_{rr}(k) \boldsymbol{\omega}_t(k), \quad (37)$$

and the cumulative log expected excess portfolio return is given by:

$$\text{E}_t \left(r_{p,t+k}^{(k)} - r_f^{(k)} k \right) + \frac{1}{2} \text{Var}_t \left(r_{p,t+k}^{(k)} - kr_f^{(k)} \right) = \boldsymbol{\omega}'_t \left(\text{E}_t \left(\mathbf{r}_{t+k}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right). \quad (38)$$

Note that we already have expressions for the moments in the RHS of these equations based on our VAR specification.

The Lagrangian for this problem is

$$\mathcal{L} = \frac{1}{2} \frac{\text{Var}_t \left(r_{p,t+k}^{(k)} - kr_f^{(k)} \right)}{k} + \lambda_f \left(\mu_p - r_f^{(k)} - \frac{1}{k} \left(\text{E}_t \left(r_{p,t+k}^{(k)} - r_f^{(k)} k \right) + \frac{1}{2} \text{Var}_t \left(r_{p,t+k}^{(k)} - kr_f^{(k)} \right) \right) \right),$$

where λ_f is the Lagrange multiplier on the expected portfolio return constraint. Note that we have already imposed the constraint that portfolio weights must add up to one in the way we define the log return on the portfolio.

Myopic portfolio rule The solution to the mean-variance program (35) is

$$\boldsymbol{\omega}_t(k) = \lambda_f \boldsymbol{\Sigma}_{rr}^{-1}(k) \left[\mathbf{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right]. \quad (39)$$

Equation for mean-variance frontier Equation (39) implies the following relation between $\sigma_p^2(k) \equiv \text{Var}_t \left(r_{p,t+k}^{(k)} - r_f^{(k)} k \right)$ and μ_p :

$$\frac{\sigma_p^2(k)}{k} = \frac{1}{A(k)} e\mu_p^2, \quad (40)$$

where $e\mu_p$ is the expected excess return on the portfolio over the riskfree asset,

$$e\mu_p = \mu_p - r_f^{(k)},$$

and

$$A(k) = \frac{1}{k} \left(\mathbf{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right)' \boldsymbol{\Sigma}_{rr}^{-1}(k) \left(\mathbf{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right).$$

Equation (40) implies that, in a standard deviation-expected return space, the frontier is a line that intersects the expected return axis at $r_f^{(k)}$, and whose slope is the Sharpe ratio.

Tangency portfolio We can now derive the weights for the tangency portfolio between frontiers (34) and (40). This is a portfolio whose weights $\boldsymbol{\omega}_{T,t}(k)$ satisfy $\boldsymbol{\omega}_{T,t}(k)' \boldsymbol{\iota} = \mathbf{1}$. That is, it is a portfolio with no loading on the risk-free asset. We can easily find the weights in this portfolio by finding the Lagrange multiplier λ_f for which $\boldsymbol{\omega}_{T,t}(k)$ in (39) satisfies this constraint:

$$1 = \boldsymbol{\omega}_{T,t}(k)' \boldsymbol{\iota} = \lambda_f \left[\mathbf{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right] (\boldsymbol{\Sigma}_{rr}^{-1}(k))' \boldsymbol{\iota}$$

or

$$\lambda_f = \frac{1}{\left[\mathbf{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right]' (\boldsymbol{\Sigma}_{rr}^{-1}(k))' \boldsymbol{\iota}},$$

so that

$$\omega_{T,t}(k) = \frac{1}{\left[\mathbb{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right]' (\boldsymbol{\Sigma}_{rr}^{-1}(k))' \boldsymbol{\iota}} \boldsymbol{\Sigma}_{rr}^{-1}(k) \left[\mathbb{E}_t \left(\mathbf{r}_{t+1}^{(k)} - r_f^{(k)} k \boldsymbol{\iota} \right) + \frac{1}{2} \boldsymbol{\sigma}_r^2(k) \right]. \quad (41)$$

Once again, it is straightforward to show that the Sharpe ratio of the tangency portfolio equals $\sqrt{A(k)}$, i.e., that the Sharpe ratio is the slope of the mean-standard deviation frontier with a riskless asset. To see this, substitute (41) into equations (37) and (38) to compute expressions for the variance and expected excess return on the tangency portfolio, and compute the ratio of the expected excess return to the standard deviation.

Table 1: Mean and standard deviation of returns and return forecasting variables, 1952.Q2-2002.Q4.

(All variables except the log dividend yield are in annualized percentage units.)

1	Mean 90-day T-bill real rate	1.518
2	Standard deviation of 90-day T-bill real rate	1.354
3	Mean excess return on stocks	6.311
4	Standard deviation of excess return on stocks	16.923
5	Sharpe ratio of stocks = (3)/(4)	0.373
6	Mean excess return on 5-year Treas. bonds	1.371
7	Standard deviation of excess return on 5-year Treas. bonds	5.649
8	Sharpe ratio of 5-year Treas. bonds = (6)/(7)	0.243
9	Mean 90-day T-bill nominal rate	5.372
10	Standard deviation of 90-day T-bill nominal rate	1.411
11	Mean log dividend yield	-3.473
12	Standard deviation of log dividend yields	0.372
13	Mean percentage dividend yield	3.323
14	Mean yield spread [†]	1.027
15	Standard deviation of yield spreads [†]	0.517

[†] The yield spread is the difference between the yield on a 5-year zero-coupon bond and the yield on a 90-day T-bill.

Table 2: VAR estimation results, 1952.Q2-2002.Q4.**A. Slopes (t-statistics in parenthesis)**

	Coefficients on lagged variables						R-squared
	(1)	(2)	(3)	(4)	(5)	(6)	
(1) log real T-bill rate	0.4029 (5.732)	0.0009 (0.167)	0.0002 (0.008)	0.2597 (3.493)	-0.0013 (-1.211)	0.3999 (2.254)	0.299
(2) log stock excess returns	1.0536 (0.977)	0.0268 (0.380)	0.3796 (1.583)	-2.1044 (-2.431)	0.0549 (3.021)	-0.1348 (-0.056)	0.095
(3) log bond excess returns	-0.1088 (-0.374)	-0.0568 (-2.672)	-0.0686 (-0.583)	0.3706 (0.944)	-0.0039 (-0.769)	3.1252 (3.169)	0.097
(4) log nominal yield on T-bills	-0.0023 (-0.074)	0.0050 (2.315)	0.0027 (0.217)	0.9539 (22.262)	0.0003 (0.772)	0.0961 (1.084)	0.872
(5) log dividend yield	-1.2428 (-1.129)	-0.0157 (-0.217)	-0.3753 (-1.489)	1.3617 (1.518)	0.9613 (52.792)	-0.3787 (-0.154)	0.951
(6) log yield spread	0.0041 (0.174)	-0.0015 (-0.869)	0.0029 (0.378)	0.0217 (0.848)	-0.0001 (-0.352)	0.7596 (13.160)	0.561


B. Quarterly percent standard deviations of residuals (diagonal) and cross-correlations of residuals (off-diagonal)

	(1)	(2)	(3)	(4)	(5)	(6)
(1) log real T-bill rate	0.567	0.221	0.357	-0.349	-0.223	0.143
(2) log stock excess returns	0.221	8.064	0.119	-0.101	-0.983	0.015
(3) log bond excess returns	0.357	0.119	2.689	-0.755	-0.141	0.147
(4) log nominal yield on T-bills	-0.349	-0.101	-0.755	0.252	0.132	-0.754
(5) log dividend yield	-0.223	-0.983	-0.141	0.132	8.175	-0.042
(6) log yield spread	0.143	0.015	0.147	-0.754	-0.042	0.172

Table 3: Tangency portfolio implied by VAR estimation and risk-free rate equal to mean real T-bill rate.

	Horizon			
	1 year (4 quarters)	10 years (40 quarters)	25 years (100 quarters)	100 years (400 quarters)
Risk free rate	1.518%	1.518%	1.518%	1.518%
Sharpe ratio	0.201	0.303	0.407	0.443
Portfolio weights				
Stocks	47.40%	145.66%	71.76%	67.10%
5-year bond	52.60%	-45.66%	28.24%	32.90%
Sum	100.00%	100.00%	100.00%	100.00%

Figure 1. Annualized Percent Standard Deviations of Real Returns Implied by Quarterly VAR(1) Estimates (1952.Q-2002.Q4)



**Figure 2. Correlations of Real Returns Implied by Quarterly VAR(1) Estimates
(1952.Q-2002.Q4)**




Figure 3. Risk of T-bills and Risk of the Global Minimum Variance Portfolio (GMV).




Figure 4. Composition of Global Minimum Variance Portfolio

