THE JOURNAL OF FINANCE • VOL. LXI, NO. 5 • OCTOBER 2006

Dynamic Portfolio Selection by Augmenting the Asset Space

MICHAEL W. BRANDT and PEDRO SANTA-CLARA*

ABSTRACT

We present a novel approach to dynamic portfolio selection that is as easy to implement as the static Markowitz paradigm. We expand the set of assets to include mechanically managed portfolios and optimize statically in this extended asset space. We consider "conditional" portfolios, which invest in each asset an amount proportional to conditioning variables, and "timing" portfolios, which invest in each asset for a single period and in the risk-free asset for all other periods. The static choice of these managed portfolios represents a dynamic strategy that closely approximates the optimal dynamic strategy for horizons up to 5 years.

SEVERAL STUDIES POINT OUT THE IMPORTANCE of dynamic trading strategies to exploit the predictability of the first and second moments of asset returns and hedge changes in the investment opportunity set. However, computing these optimal dynamic investment strategies has proven to be a rather formidable problem because closed-form solutions are only available for a few cases. While researchers have explored a variety of numerical solution methods, including solving partial differential equations, discretizing the state-space, and using Monte Carlo simulation, these techniques are out of reach for most practitioners and thus they remain largely in the ivory tower. The workhorse of portfolio optimization in industry remains the static Markowitz approach.

Our paper presents a novel approach to dynamic portfolio selection that is no more difficult to implement than the static Markowitz model. The idea is to expand the asset space to include simple (mechanically) managed portfolios and compute the optimal *static* portfolio within this extended asset space. The intuition is that a static choice of managed portfolios is equivalent to a dynamic strategy. One can therefore approximate the optimal dynamic strategy by a fixed combination of mechanically managed portfolios. We consider managed portfolios of two types, namely "conditional" and "timing" portfolios. Conditional managed portfolios are constructed along the lines of Hansen and

*Brandt is at the Fuqua School of Business, Duke University, and is also affiliated with the NBER; Santa-Clara is at the Anderson School of Management, University of California Los Angeles, and is also affiliated with the NBER. We thank Michael Brennan, John Campbell, Rob Engle, Rick Green, Francis Longstaff, Eduardo Schwartz, Rob Stambaugh, Rossen Valkanov, Luis Viceira, Shu Yan, an anonymous referee, and seminar participants at the 2005 American Finance Association meetings, ABP Investments, Baruch College, Barclays Global Investors, European Central Bank, Lehman Brothers, New York University, and Morgan Stanley for helpful comments.

Richard (1987).[1] Specifically, for each variable that affects the distribution of returns and for each basis asset, we consider a portfolio that invests in the basis asset an amount that is proportional to the level of the conditioning variable. Timing portfolios invest in each asset for a single period and in the risk-free rate for all other periods. Timing portfolios mimic strategies that buy and sell the asset over time. For example, holding a constant amount of all the timing portfolios that are related to a single asset approximates a strategy that holds a constant proportion of wealth in the asset. In contrast, hedging demands induce the investor to hold different amounts of the timing portfolios across time.

Given an expanded asset space with managed portfolios, we can use the Markowitz solution to find the optimal strategy for a mean–variance investor. The optimal strategy is a combination of managed portfolios. However, it is trivial to recover the corresponding investment in the basis assets at each point in time given the values of the conditioning variables. The weight invested in each basis asset at each point in time is a simple linear function of the state variables. Thus, our approach is equivalent to parameterizing the portfolio policy as a function of the state variables and then maximizing the investor's utility by choosing optimally the coefficients of this function.

The advantage of framing the dynamic portfolio problem in a static context is that all the refinements developed over the years for the Markowitz model become available. These include the use of portfolio constraints, shrinkage estimation, and the combination of an investor's prior beliefs with the information contained in the history of returns.

In general, our approach relies on sample moments of the long-horizon returns of the expanded set of assets. However, given the finite size of the sample of returns, we cannot address problems with very long horizon by simply expanding the asset space. For example, the Ibbotson database covers roughly 80 years of data, which corresponds to only eight nonoverlapping 10-year returns and renders our approach subject to small-sample problems for problems with such long horizons. For these cases, it is better to use a model of the dynamics of returns and state variables. We show that if the log returns on the basis assets and the log state variables follow a vector auto-regression (VAR) with normally distributed innovations, as is typically assumed in the line of research initiated by Campbell and Viceira (1999),[2] the long-horizon moments of returns can all be expressed in terms of the parameters of the VAR. In this special but popular case, we use our approach to obtain approximate closed-form solutions for the finite-horizon dynamic portfolio choice that complements the approximate closed-form solutions for the infinite-horizon case with intermediate consumption derived by Campbell and Viceira (1999).

Our approach is similar in spirit to that of Cox and Huang (1989) and its empirical implementation by Aït-Sahalia and Brandt (2005). In these papers, the dynamic portfolio problem is solved in two steps. The investor first chooses

[1] Hansen and Richard (1987) introduced this idea to develop tests of conditional asset pricing models. Bansal and Harvey (1996) use conditional portfolios in performance evaluation.

[2] See also the survey by Campbell and Viceira (2002).

the optimal portfolio of Arrow–Debreu securities and then figures out how to replicate this portfolio by dynamically trading the basis assets or derivatives on the basis assets. In contrast, in this paper, we solve the portfolio problem in one step as the optimal choice across simple dynamic trading strategies. Note also that while the Cox–Huang approach requires that financial markets be complete, for only then all Arrow–Debreu securities can be replicated, we do not need to assume market completeness since the investor only chooses among feasible strategies. On the other hand, Cox and Huang (1989) provide the exact solution to the portfolio problem, whereas our approach offers only an approximation.

Our paper also relates to Ferson and Siegel (2001). Assuming that the conditional mean vector and covariance matrix of asset returns are known functions of the state variables, they derive the optimal portfolio weights by maximizing a mean–variance utility function (in an unconditional sense similar to ours). They then show that the resulting portfolio weights are also functions of the state variables since the weights depend on the conditional means, variances, and covariances of asset returns. In contrast, we model the portfolio weights directly as functions of the state variables, and we find the coefficients of these functions that maximize the investor's utility. Our portfolio weights implicitly take into account the impact of the state variables on the means, variances, and covariances of asset returns since all of these moments affect the portfolio's expected return and risk, and in turn the investor's expected utility. Thus, our method can be interpreted as an approximation of the solution offered by Ferson and Siegel. For instance, by postulating that the optimal portfolio weights are linear in the state variables, we implicitly constrain the forms of the mean vector and the covariance matrix of returns as functions of the state variables.

However, the two methods are quite different when applied in practice. To use Ferson and Siegel's approach, we need to estimate conditional means, variances, and covariances of returns as functions of the state variables. While one can easily estimate conditional mean functions by regressing returns on the state variables, it is notoriously difficult to estimate a conditional covariance matrix as a function of state variables in a manner that guarantees positive semidefiniteness at all times. In contrast, estimating the portfolio weight function in our approach does not require imposing any sort of nonlinear constraints. Furthermore, our approach has the advantage of being much more parsimonious. Suppose we are interested in forming optimal portfolios of N assets. With Ferson and Siegel's approach, we have to estimate N functions of the state variables for the expected return vector and $N(N + 1)/2$ functions for the covariance matrix. With our approach, we only need to estimate N functions for the optimal portfolio weights. The gains in computation and estimation precision are evident.

A few words of caution are in order. Given that our approximation to the multiperiod optimization problem ignores compounding of returns, our method cannot address portfolio choice problems with very long horizons. The loss in certainly equivalent from using our approximation as opposed to an exact solution can exceed 10% at a 10-year horizon and 25% at a 20-year horizon.

Furthermore, our approach cannot handle endogenous state variables obtained from an investor's previous decisions. Endogenous state variables are important in problems with transaction costs, taxes, and other frictions. Finally, our method is only applicable to preferences over final wealth. It cannot be applied to problems with intermediate consumption.

The paper proceeds as follows. First, we describe our approach in Sections I.A and I.B. We then illustrate the mechanics of our approach through a simple example in Section I.C, and we examine its accuracy in Section I.D. Section II deals with the special case in which the log returns and log state variables follow a Gaussian VAR, and Section III discusses how several refinements of the static Markowitz approach can be directly applied to our approach. We illustrate our approach through an empirical application in Section IV, and in Section V we conclude.

I. The Method

We solve a conditional portfolio choice problem with parametrized portfolio weights of the form $x_t = \theta z_t$, where z_t denotes a vector of state variables and θ is a matrix of coefficients. This conditional portfolio choice problem is mathematically equivalent to solving an unconditional problem within an augmented asset space that includes naively managed zero-investment portfolios with excess returns of the form z_t times the excess return of each basis asset. In subsection A we establish this idea in the context of a single-period problem, and in subsection B we extend the approach to the multiperiod case. We then illustrate both cases in a simple example (subsection C), and finally we examine the accuracy of the solutions in a numerical experiment (subsection D).

A. Single-Period Problem

We consider the problem of an investor who maximizes the conditional expected value of a quadratic utility function over next period's wealth, W_{t+1}

$$\max \mathrm{E}_t \left[W_{t+1} - \frac{b_t}{2} W_{t+1}^2 \right], \tag{1}$$

where b_t is positive and sufficiently small to ensure that the marginal utility of wealth remains positive. Let R_t^f be the gross risk-free rate and $r_{t+1}^p = R_{t+1}^p - R_t^f$ be the excess return on the investor's portfolio from t to $t+1$. (Throughout the paper we use capital letters to denote gross returns and lower-case letters to denote excess returns. We date all variables with a subscript that corresponds to the time at which the variable is known. For example, returns of risky assets from time t to time $t+1$ are denoted by R_{t+1}. The risk-free rate for the same period is denoted by R_t^f, since it is known at the beginning of the return period.) Given this notation, we have

$$W_{t+1} = W_t \left(R_t^f + r_{t+1}^p \right). \tag{2}$$

Substituting equation (2) into equation (1) and simplifying, we obtain

$$\max \mathrm{E}_t \left[cte + r_{t+1}^p - \frac{b_t W_t}{2\left(1 - b_t W_t R_t^f\right)} \left(r_{t+1}^p\right)^2 \right], \tag{3}$$

where *cte* contains terms that are constant given the information available at time t. For simplicity, we rewrite the problem as

$$\max \mathrm{E}_t \left[r_{t+1}^p - \frac{\gamma}{2} \left(r_{t+1}^p\right)^2 \right], \tag{4}$$

ignoring the constant term and letting γ represent a positive constant.

Denote the vector of portfolio weights on the risky assets at time t by x_t. The above optimization problem then becomes

$$\max_{x_t} \mathrm{E}_t \left[x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top r_{t+1} r_{t+1}^\top x_t \right], \tag{5}$$

where $r_{t+1} = R_{t+1} - R_t^f$ is the vector of excess returns on the N risky assets. By formulating the problem in terms of excess returns, we are implicitly assuming that the remainder of the portfolio's value is invested in the risk-free asset with return R_t^f.

When the returns are independent and identically distributed (i.i.d.) and the portfolio weights are constant over time, that is $x_t = x$, we can replace the conditional expectation with an unconditional expectation and solve for the weights according to

$$x = \frac{1}{\gamma} \mathrm{E}\left[r_{t+1} r_{t+1}^\top\right]^{-1} \mathrm{E}[r_{t+1}]. \tag{6}$$

This is the well-known Markowitz solution, which can be implemented in practice by replacing the population moments by sample averages

$$x = \frac{1}{\gamma} \left[\sum_{t=1}^{T-1} r_{t+1} r_{t+1}^\top \right]^{-1} \left[\sum_{t=1}^{T-1} r_{t+1} \right]. \tag{7}$$

(Note that the $1/T$ terms in the sample averages cancel.)

Consider now the more realistic case of returns that are not i.i.d. and assume that the optimal portfolio policies are linear in a vector of K state variables (the first of which we will generally take to be a constant). Then

$$x_t = \theta z_t, \tag{8}$$

where θ is an $N \times K$ matrix. Our problem is now

$$\max_{\theta} \mathrm{E}_t \left[(\theta z_t)^\top r_{t+1} - \frac{\gamma}{2} (\theta z_t)^\top r_{t+1} r_{t+1}^\top (\theta z_t) \right]. \tag{9}$$

From linear algebra we can use the result that

$$(\theta z_t)^\top r_{t+1} = z_t^\top \theta^\top r_{t+1} = \mathrm{vec}(\theta)^\top (z_t \otimes r_{t+1}), \tag{10}$$

where vec(θ) piles up the columns of matrix θ into a vector and $\otimes$ is the Kronecker product of two matrices, and can write

$$\tilde{x} = \mathrm{vec}(\theta) \tag{11}$$

$$\tilde{r}_{t+1} = z_t \otimes r_{t+1}. \tag{12}$$

Our problem can now be written as

$$\max_{\tilde{x}} \mathrm{E}_t \left[\tilde{x}^\top \tilde{r}_{t+1} - \frac{\gamma}{2} \tilde{x}^\top \tilde{r}_{t+1} \tilde{r}_{t+1}^\top \tilde{x} \right]. \tag{13}$$

Since the same $\tilde{x}$ maximizes the conditional expected utility at all dates t, it also maximizes the unconditional expected utility

$$\max_{\tilde{x}} \mathrm{E} \left[\tilde{x}^\top \tilde{r}_{t+1} - \frac{\gamma}{2} \tilde{x}^\top \tilde{r}_{t+1} \tilde{r}_{t+1}^\top \tilde{x} \right], \tag{14}$$

which corresponds to the problem of finding the unconditional portfolio weights $\tilde{x}$ for the expanded set of $(N \times K)$ assets with returns $\tilde{r}_{t+1}$. The expanded set of assets can be interpreted as managed portfolios, each of which invests in a single basis asset an amount that is proportional to the value of one of the state variables. We term these "conditional portfolios."

It follows that the optimal $\tilde{x}$ is

$$\begin{aligned} \tilde{x} &= \frac{1}{\gamma} \mathrm{E}\big[\tilde{r}_{t+1} \tilde{r}_{t+1}^\top\big]^{-1} \mathrm{E}[\tilde{r}_{t+1}] \\ &= \frac{1}{\gamma} \mathrm{E}\big[\big(z_t z_t^\top\big) \otimes \big(r_{t+1} r_{t+1}^\top\big)\big]^{-1} \mathrm{E}[z_t \otimes r_{t+1}], \end{aligned} \tag{15}$$

which we can again implement in practice by replacing the population moments by sample averages

$$\tilde{x} = \frac{1}{\gamma} \left[\sum_{t=0}^{T} \big(z_t z_t^\top\big) \otimes \big(r_{t+1} r_{t+1}^\top\big) \right]^{-1} \left[\sum_{t=0}^{T} z_t \otimes r_{t+1} \right]. \tag{16}$$

From this solution we can trivially recover the weight invested in each of the basis assets by adding the corresponding products of elements of $\tilde{x}$ and z_t.

Note that the solution (16) depends only on the data and hence does not require any assumptions about the distribution of returns apart from stationarity. In particular, the solution does not require any assumptions about how the distribution of returns depends on the state variables. Thus, the state variables can predict time variation in the first-, second-, and, if we consider more general utility functions, even higher-order moments of returns. As Brandt (1999) and Aït-Sahalia and Brandt (2001) emphasize, the advantage of focusing directly on the portfolio weights is that we bypass estimation of the conditional return distribution. This intermediate estimation step typically involves ad hoc distributional assumptions and inevitably misspecified models for the conditional moments of returns. In contrast, estimating the conditional portfolio weights in a single step is robust to misspecification of the conditional return distribution. It can also result in more precise estimates if the dependence of the optimal portfolio weights on the state variables is less noisy than the dependence of the return moments on the state variables.

At this point, it is instructive to compare our approach to that of Ferson and Siegel (2001). They assume that the conditional expected returns and the conditional variances and covariances of asset returns are known functions of the state variables, that is,

$$r_{t+1} = \mu(z_t) + \epsilon_{t+1}, \tag{17}$$

where the conditional covariance matrix of ϵ_{t+1} is $\Sigma(z_t)$. Ferson and Siegel then derive the mean–variance optimal portfolio weights as a function of the state variables

$$x(z_t) = \pi\left(\mu(z_t) - R^f\iota\right)^\top \Lambda(z_t), \tag{18}$$

where

$$\Lambda(z_t) = \left[\left(\mu(z_t) - R^f\iota\right)\left(\mu(z_t) - R^f\iota\right)^\top + \Sigma(z_t)\right]^{-1}, \tag{19}$$

ι is a vector of ones, and π is a constant.

Our approach of modeling the portfolio weights as a function of the state variables can be seen as an approximation of the solution provided by Ferson and Siegel. For instance, postulating that the portfolio weights are linear in the state variables,

$$x(z_t) = \theta z_t, \tag{20}$$

implicitly constrains the functional forms of $\mu(z)$ and $\Sigma(z)$ in equations (18) and (19). Ferson and Siegel show that when the returns are homoskedastic, the optimal portfolio weights are approximately linear in the expected returns for an extended range of the state variables around their unconditional means. Therefore, if the expected returns are linear in the state variables, the portfolio weights will also be linear in the state variables. Of course, homoskedastic returns with linear means is only one of many return models that deliver approximately linear portfolio weights. Also, our approach can easily accommodate nonlinear portfolio weights by simply including nonlinear transformations of the state variables in the portfolio weight functions.

In practice, applying Ferson and Siegel's approach raises a number of issues. While one can easily estimate the conditional mean functions $\mu(z_t)$ by regressing excess returns r_{t+1} on the state variables z_t, it is notoriously difficult to estimate a conditional covariance matrix $\Sigma(z_t)$ as a function of the state variables in a manner that guarantees positive semidefiniteness at all times. Estimating the portfolio weight function in our approach does not require imposing any sort of nonlinear constraints. Furthermore, our approach has the advantage of being much more parsimonious. Suppose that we are interested in forming optimal portfolios of N assets. With Ferson and Siegel's approach, we have to estimate

N functions of the state variables for the expected return vector and $N(N+1)/2$ functions for the covariance matrix, whereas with our approach, we only need to estimate N functions for the optimal portfolio weights. The gains in computation and estimation precision are evident.

Since we express the portfolio problem in an estimation context, we can use standard sampling theory to compute standard errors for the portfolio weights and then test hypotheses about them. Specifically, following Britten-Jones (1999), we can interpret the solution (16) as being proportional (with constant of proportionality $1/\gamma$) to the coefficients of a standard ordinary least squares (OLS) regression of a vector of ones on the excess returns $\tilde{r}_{t+1}$. This allows us to compute standard errors for $\tilde{x}$ from the standard errors of the regression coefficients. These standard errors can be used to test, for example, whether some state variable is a significant determinant of the portfolio policy. Using our notation, the covariance matrix of the vector $\tilde{x}$ is

$$\frac{1}{\gamma^2}\frac{1}{T-N\times K}(\iota_T-\tilde{r}\tilde{x})^\top(\iota-\tilde{r}\tilde{x})(\tilde{r}^\top\tilde{r})^{-1}, \tag{21}$$

where ι_T denotes a $T\times 1$ vector of ones and $\tilde{r}$ is a $T\times K$ matrix with the time series of returns of the K managed portfolios.

As we already mentioned, the assumption that the optimal portfolio weights are linear functions of the state variables is innocuous because z_t can include nonlinear transformations of a set of more basic state variables y_t. This means that the linear portfolio weights can be interpreted as a more general portfolio policy function $x_t = g(y_t)$ for any $g(\cdot)$ that can be spanned by a polynomial expansion in the more basic state variables y_t. In other words, in principle our approach can accommodate very general dependence of the optimal portfolio weights on the state variables.

In practice, we need to choose a finite set of state variables and possible nonlinear transformations of these state variables to include in the portfolio policy. From a statistical perspective, variable selection for modeling portfolio weights is no different from variable selection for modeling returns. Variables can be chosen on the basis of individual t-tests and joint F-tests computed with the covariance matrix of the portfolio weights in equation (21), or on the basis of out-of-sample performance. From an economic perspective, however, there are distinct advantages to focusing directly on the optimal portfolio weights. As Aït-Sahalia and Brandt (2001) demonstrate, it is more natural in an asset allocation framework to choose variables that predict optimal portfolio weights than it is to choose variables that predict return moments. In particular, a variable may be a statistically important predictor of both means and variances but be useless for determining optimal portfolio weights because the variation in the moments offset each other (e.g., the corresponding conditional Sharpe ratio is small).

Finally, we can extend our approach to allow some or all of the state variables to be asset-specific. In a companion paper, Brandt, Santa-Clara, and Valkanov (2005), we study optimal stock portfolios by parameterizing the weight invested in each stock as a function of the company's characteristics, including its book-to-market ratio, market capitalization, and return over the past year. Importantly, the parameters of the weight function are constrained to be the same for all stocks, which makes the problem highly tractable and computationally efficient. The resulting optimal portfolios (of this very large set of assets) do not suffer from exploding weights (as mean–variance efficient portfolios often do), and they deliver outstanding performance both in and out of sample.

B. Multiperiod Problem

The idea of augmenting the asset space with naively managed portfolios extends to the multiperiod case. Consider an investor who maximizes the two-period mean-variance objective

$$\max \mathrm{E}_t \left[r^p_{t\to t+2} - \frac{\gamma}{2}\left(r^p_{t\to t+2}\right)^2 \right], \tag{22}$$

where $r^p_{t\to t+2}$ denotes the excess return of the two-period investment strategy

$$\begin{aligned} r^p_{t\to t+2} &= \left(R^f_t + x_t^\top r_{t+1}\right)\left(R^f_{t+1} + x_{t+1}^\top r_{t+2}\right) - R^f_t R^f_{t+1} \\ &= x_t^\top\left(R^f_{t+1} r_{t+1}\right) + x_{t+1}^\top\left(R^f_t r_{t+2}\right) + \left(x_t^\top r_{t+1}\right)\left(x_{t+1}^\top r_{t+2}\right). \end{aligned} \tag{23}$$

The first line of this expression shows why we refer to $r^p_{t\to t+2}$ as a two-period excess return. The investor borrows a dollar at date t and allocates it to the risky and risk-free assets according to the first-period portfolio weights x_t. After the first period, at date $t+1$, the one-dollar investment results in $(R^f_t + x_t^\top r_{t+1})$ dollars, which the investor then allocates again to the risky and risk-free assets according to the second-period portfolio weights x_{t+1}. Finally, at date $t+2$, the investor has $(R^f_t + x_t^\top r_{t+1})(R^f_{t+1} + x_{t+1}^\top r_{t+2})$ dollars but must pay $R^f_t R^f_{t+1}$ dollars for the principal and interest of the one-dollar loan. The remainder is the two-period excess return.

The second line of equation (23) decomposes the two-period excess return into three terms. The first two terms have a natural interpretation as the excess return of investing in the risk-free rate in the first (second) period and in the risky asset in the second (first) period. To see that $x_t^\top(R^f_{t+1} r_{t+1})$ is a two-period excess return from investing in risky assets in the first period and the risk-free asset in the second period, just follow the argument above letting $x_{t+1} = 0$. Investing the first-period proceeds of $(R^f_t + x_t^\top r_{t+1})$ in the risk-free asset in the second period yields $(R^f_t + x_t^\top r_{t+1})R^f_{t+1}$. After paying back $R^f_t R^f_{t+1}$, the investor is left with an excess return of $x_t^\top(R^f_t r_{t+1})$. Note that the portfolio weights on these two intertemporal portfolios are the same as the weights on the risky asset in the first and second periods, respectively. The third term in this expression captures the effect of compounding.

Comparing the first two terms to the third, we see that the latter is two orders of magnitude smaller than the former. The return $(x_t^\top r_{t+1})(x_{t+1}^\top r_{t+2})$ is a product of two single-period excess returns, which means that its magnitude is typically of the order of 1/100$^{\text{th}}$ of a percent per year. The returns on the first two portfolios, in contrast, are products of a gross return (R_t^f or R_{t+1}^f) and an excess return (r_{t+1} or r_{t+2}), so their magnitudes are likely to be percent per year.

Given that the compounding term is orders of magnitude smaller than the two intertemporal portfolios, we ignore it for now. (We discuss the effect of ignoring the compounding term below.) The two-period portfolio choice is then simply a choice between two intertemporal portfolios, one that holds the risky asset in the first period only and the other that holds the risky asset in the second period only. We refer to these two portfolios as "timing portfolios." We can then solve the dynamic problem as a simple static choice between these two managed portfolios. In particular, for the two-period case, the sample analogue of the optimal portfolio weights is given by

$$\tilde{x} = \frac{1}{\gamma}\left[\sum_{t=1}^{T-2} \tilde{r}_{t\to t+2}\tilde{r}_{t\to t+2}^{\top}\right]^{-1}\left[\sum_{t=1}^{T-2} \tilde{r}_{t\to t+2}\right], \tag{24}$$

where $\tilde{r}_{t\to t+2} = [R_{t+1}^f r_{t+1}, R_t^f r_{t+2}]$. The first set of elements of $\tilde{x}$ (corresponding to the returns $R_{t+1}^f r_{t+1}$) represents the fraction of wealth invested in the risky assets in the first period, and the second set of elements (corresponding to $R_t^f r_{t+2}$) represents the fraction of wealth invested in the risky assets in the second period.

In a general H-period problem, we proceed in exactly the same fashion. We construct a set of timing portfolios

$$\tilde{r}_{t\to t+H} = \left\{\prod_{\substack{i=0\\ i\neq j}}^{H-1} R_{t+i}^f r_{t+j+1}\right\}_{j=0}^{H-1}, \tag{25}$$

where each term represents a portfolio that invests in risky assets in period $t+j$ and in the risk-free rate in all other periods $t+i$, with $i \neq j$. Again, the sample analogue of the optimal portfolio weights is then given by the static solution

$$\tilde{x} = \frac{1}{\gamma}\left[\sum_{t=1}^{T-H} \tilde{r}_{t\to t+H}\tilde{r}_{t\to t+H}^{\top}\right]^{-1}\left[\sum_{t=1}^{T-H} \tilde{r}_{t\to t+H}\right]. \tag{26}$$

It is important to realize that, in contrast to a long-horizon buy-and-hold problem, the random components of the timing portfolios are nonoverlapping. We therefore avoid the usual statistical problems associated with overlapping long-horizon returns. Note, however, that as the length of the horizon H increases, we lose observations for computing the mean and covariance matrix of $\tilde{r}_{t\to t+H}$, which may compromise the statistical precision of the solution.

We can naturally combine the ideas of conditional portfolios and timing portfolios by replacing the risky returns r_{t+j+1} in equation (25) with the conditional portfolio returns $z_{t+j} \otimes r_{t+j+1}$. The resulting optimal portfolio weights $\tilde{x}$ from equation (26) then provide the optimal allocations to the conditional portfolios at each time $t+j$.

The obvious appeal of our approach is its simplicity. Of course, this simplicity comes with drawbacks. First, by ignoring the compounding terms, our approach no longer provides the exact solution to the multiperiod problem. Writing out the return to an H-period dynamic portfolio strategy analogous to the two-period case in equation (23) shows that the multiperiod portfolio returns are only spanned when we include the compounding terms in the static portfolio problem. Unfortunately, the presence of the compounding terms imposes a set of nonlinear constraints on the static portfolio weights. The portfolio weights on the compounding terms are constrained to be products of the portfolio weights on the timing portfolios. Due to the nonlinearity of these constraints, solving the static constrained problem with compounding terms for a large number of assets and/or a large number of rebalancing periods is not much simpler than solving the corresponding dynamic problem using numerical optimization techniques. Our suggestion is to ignore the compounding terms on the grounds that they are orders of magnitude smaller than the timing portfolio returns. However, in ignoring the compounding terms, our solution is at best a good approximation of the solution to the multiperiod problem. The quality of the approximation is naturally specific to each application. Intuitively, it depends on the growth rate of wealth per period and on the number of periods considered.

To better understand this issue consider the exact H-period portfolio excess return

$$\begin{aligned} r^p_{t\to t+H} &= \sum_{j=0}^{H-1} x_{t+j}^\top \prod_{\substack{i=0\\ i\neq j}}^{H-1} R^f_{t+i} r_{t+j+1} \\ &\quad + \sum_{j=0}^{H-1} \sum_{\substack{k=0\\ k\neq j}}^{H-1} x_{t+j}^\top \prod_{\substack{i=0\\ i\neq j}}^{H-1} R^f_{t+i} r_{t+j+1} x_{t+k}^\top \prod_{\substack{i=0\\ i\neq k}}^{H-1} R^f_{t+i} r_{t+k+1} \dots \qquad (27) \\ &\quad + \prod_{j=0}^{H-1} x_{t+j}^\top r_{t+j+1}. \end{aligned}$$

Our approximation uses the first H terms in this expression (in the first line of the equation) and disregards $2^H - H - 1$ terms of lower magnitude. It is important to note that the number of terms we disregard grows exponentially with the number of rebalancing periods H and the importance of each of these terms depends on the magnitude of the one-period returns. Thus, the quality of the approximation is likely to deteriorate both with the horizon of the portfolio problem and with the rebalancing frequency. In subsection D we use a simulation experiment to provide some evidence on the quality of the approximation for different horizons and rebalancing frequencies.

The second drawback of our approach is that it can be quite data intensive for problems with very long horizons. For example, suppose we want to solve a 10-year portfolio choice problem with quarterly rebalancing using a 60-year post-war sample of quarterly returns and state variable realizations. Since each timing portfolio involves a 10-year return, we would only have six independent observations to compute the moments of the timing portfolio returns and hence the optimal portfolio weights. The obvious way to overcome this data issue is to impose a statistical model for the returns and state variables that allows us to compute the long-horizon moments analytically (or by simulation) from the parameters of the statistical model. Specifically, if the log returns on the basis assets and the log state variables follow a VAR with normally distributed innovations, the long-horizon moments can be expressed in terms of the parameters of the VAR. This use of a statistical model allows us to solve dynamic portfolio choice problems with arbitrarily long horizons using only a finite data sample. We elaborate on this idea in Section II.

C. Illustrative Example

To illustrate more concretely the mechanics of our approach, consider a time series of only six observations (for simplicity) of excess returns for two risky assets, namely, a stock denoted by s and a bond denoted by b

$$\begin{bmatrix} r_1^s & r_1^b \\ r_2^s & r_2^b \\ \cdots & \cdots \\ r_6^s & r_6^b \end{bmatrix}. \tag{28}$$

The optimal static portfolio in equation (7) directly gives us the weight x^s invested in the stock and the weight x^b invested in the bond (with the remainder invested in the risk-free asset). The solution takes into account the sample covariance matrix of asset returns and the vector of sample mean excess returns.

Suppose now that there is one conditioning variable, such as the dividend yield or the spread between long and short Treasury yields, which affects the conditional distribution of returns. We observe a time series of this state variable

$$\begin{bmatrix} z_0 \\ z_1 \\ \cdots \\ z_5 \end{bmatrix}, \tag{29}$$

where the dating reflects the fact that z is known at the beginning of each return period. We take the information in the conditioning variable into account by estimating a portfolio policy that depends on it. To do this, we expand the matrix of returns (28) in the following manner:

$$\begin{bmatrix} r_1^s & r_1^b & z_0 r_1^s & z_0 r_1^b \\ r_2^s & r_2^b & z_1 r_2^s & z_1 r_2^b \\ \cdots & \cdots & \cdots & \cdots \\ r_6^s & r_6^b & z_5 r_6^s & z_5 r_6^b \end{bmatrix}, \tag{30}$$

and we then compute the optimal static portfolio of this expanded set of assets. This static solution gives us a vector of four portfolio weights $\tilde{x}$ that correspond to each of the basis assets and managed portfolios in the matrix above. We find the weight invested in the stock in each period by using the first and third elements of $\tilde{x}$, that is, $x_t^s = \tilde{x}_1 + \tilde{x}_3 z_t$. Similarly, the weight invested in the bond in each period is $x_t^b = \tilde{x}_2 + \tilde{x}_4 z_t$. Note that when we use the Markowitz solution (7) on the matrix of returns of the expanded asset set (30), the covariance matrix and vector of means take into account both the covariances among returns and between returns and lagged state variables. The latter covariances capture the impact of return predictability on the optimal portfolio policy.

Consider now a two-period portfolio choice problem. We construct the matrix of returns of the timing portfolios as described in equation (25)

$$\begin{bmatrix} r_1^s R_1^f & R_0^f r_2^s & r_1^b R_1^f & R_0^f r_2^b \\ r_3^s R_3^f & R_2^f r_4^s & r_3^b R_3^f & R_2^f r_4^b \\ r_5^s R_5^f & R_4^f r_6^s & r_5^b R_5^f & R_4^f r_6^b \end{bmatrix}. \tag{31}$$

This matrix contains two-period nonoverlapping returns of four trading strategies. The corresponding optimal portfolio vector $\tilde{x}$ gives us the weights on "stock in period 1," "stock in period 2," "bond in period 1," and "bond in period 2." The covariance matrix and vector of means that show up in the static portfolio solution account for the contemporaneous covariances of returns as well as the one-period serial covariances of returns. The latter covariances induce hedging demands.

Finally, we consider a two-period problem with the conditioning variable. The returns of the expanded asset set are

$$\begin{bmatrix} r_1^s R_1^f & R_0^f r_2^s & r_1^b R_1^f & R_0^f r_2^b & z_0 r_1^s R_1^f & R_0^f z_1 r_2^s & z_0 r_1^b R_1^f & R_0^f z_1 r_2^b \\ r_3^s R_3^f & R_2^f r_4^s & r_3^b R_3^f & R_2^f r_4^b & z_2 r_3^s R_3^f & R_2^f z_3 r_4^s & z_2 r_3^b R_3^f & R_2^f z_3 r_4^b \\ r_5^s R_5^f & R_4^f r_6^s & r_5^b R_5^f & R_4^f r_6^b & z_4 r_5^s R_5^f & R_4^f z_5 r_6^s & z_4 r_5^b R_5^f & R_4^f z_5 r_6^b \end{bmatrix}. \tag{32}$$

The optimal portfolio of these eight assets now includes, for example, the allocation to "stock in period 1, conditional on the level of z." The portfolio solution takes into account the covariances between returns and state variables over subsequent periods.

D. Importance of the Compounding Terms

Our approach to the multiperiod portfolio problem relies critically on the presumption that the compounding terms (i.e., the cross-products of the excess returns in different time periods) are negligible relative to the returns on the timing portfolios. We now examine to what extent and under which circumstances this is valid.

We apply our method to the following model for monthly excess stock and bond returns (the basis assets) and the term spread (the state variable)

$$\begin{bmatrix} \ln\left(1+r^s_{t+1}\right) \\ \ln\left(1+r^b_{t+1}\right) \\ z_{t+1} \end{bmatrix} = \begin{bmatrix} 0.0059 \\ 0.0007 \\ -0.0028 \end{bmatrix} + \begin{bmatrix} 0.0060 \\ 0.0035 \\ 0.9597 \end{bmatrix} \times z_t + \begin{bmatrix} \epsilon^s_{t+1} \\ \epsilon^b_{t+1} \\ \epsilon^z_{t+1} \end{bmatrix}, \tag{33}$$

with

$$\begin{bmatrix} \epsilon^s_{t+1} \\ \epsilon^b_{t+1} \\ \epsilon^z_{t+1} \end{bmatrix} \sim \text{MVN}\left[0, \begin{bmatrix} 0.0018 & 0.0002 & -0.0005 \\ 0.0002 & 0.0006 & 0.0007 \\ -0.0005 & 0.0007 & 0.0802 \end{bmatrix}\right]. \tag{34}$$

The choice of state variable is based on our empirical results in Section IV, where we identify the term spread as an important return predictor (other important predictors include the dividend yield and the detrended short-term interest rate). The functional form of the model follows the literature on portfolio choice under predictability and is also related to our setup in Section II. The parameter values are OLS estimates based on monthly data from January 1945 through December 2000.

To assess the importance of the compounding terms in the solution of the multiperiod portfolio problem, we compare portfolio policies that ignore the compounding terms (using our simplified approach based on the timing portfolios) to policies that incorporate the compounding terms (obtained through numerical optimization). We label these solutions "approximate" and "exact," respectively.[3] Intuitively, there are two factors that affect the role of the compounding terms, specifically, the rebalancing frequency and the portfolio horizon. The less frequently the portfolio is rebalanced, the larger are the magnitudes of the excess returns per period, and therefore the larger are the magnitudes of the compounding terms; the longer the horizon, the more compounding terms there are in the expanded budget constraint. Hence, we study multiperiod portfolio problems with rebalancing frequencies ranging from monthly to annual and horizons ranging from 1 to 20 years.

The results of our experiments are displayed in Table I. The table describes the multiperiod returns from the approximate and exact portfolio policies for an investor with quadratic utility and $\gamma = 5$, the value we use in our empirical application. Panel A presents the results for unconditional portfolio policies

[3] The exact solution is obtained by numerically maximizing the expected utility of terminal wealth with respect to the portfolio weights in every period. For a given set of portfolio weights, the moments of the multiperiod portfolio returns are evaluated using 5,000,000 data points simulated from the model (33). To keep the comparison as fair as possible and to abstract from sampling error, we use the same simulations to evaluate the moments of the timing portfolios for the approximate solution.

Table I

Approximation Error in Multiperiod Portfolio Policies

This table describes the difference between the approximate and exact portfolio policies for an investor with quadratic utility and $\gamma = 5$. The approximate solution is based on timing portfolios that ignore the compounding of excess returns over time. The exact solution takes compounding into account. Panel A is for unconditional portfolio policies involving stocks and bonds. Panel B is for conditional portfolio policies in which the stock and bond returns are also scaled by the term spread. Each panel reports the average absolute difference between the approximate versus exact portfolio weights, the level and percentage (in parentheses) difference in the Sharpe ratios of the two portfolios, and the equalization fee, defined as the yearly fee the investor would be willing to pay to be able to use the exact instead of approximate allocations, expressed both as a level and a fraction of the certainty equivalent of the exact allocation (in parentheses).

Panel A: Unconditional Policies						
	Monthly Rebalancing			Quarterly Rebalancing		
	Horizon			Horizon		
	1-Year	2-Year	5-Year	2-Year	5-Year	10-Year
Avg $\|\Delta w\|$	0.0279	0.0347	0.0372	0.0297	0.0341	0.0573
Δ Sharpe ratio	0.0012 (0.21%)	0.0042 (0.53%)	0.0168 (1.39%)	0.0032 (0.41%)	0.0150 (1.25%)	0.0285 (2.94%)
Equalization fee	0.0002 (0.75%)	0.0008 (1.66%)	0.0039 (3.80%)	0.0006 (1.29%)	0.0036 (3.54%)	0.0166 (11.46%)
	Semi-Annual Rebalancing			Annual Rebalancing		
	Horizon			Horizon		
	2-Year	5-Year	10-Year	5-Year	10-Year	20-Year
Avg $\|\Delta w\|$	0.0255	0.0322	0.0386	0.0285	0.0361	0.0632
Δ Sharpe ratio	0.0022 (0.29%)	0.0131 (1.10%)	0.0407 (2.54%)	0.0099 (0.85%)	0.0340 (2.16%)	0.1407 (6.88%)
Equalization fee	0.0004 (0.86%)	0.0032 (3.16%)	0.0147 (6.82%)	0.0023 (2.30%)	0.0127 (5.95%)	0.1852 (25.58%)

Panel B: Conditional Policies						
	Monthly Rebalancing			Quarterly Rebalancing		
	Horizon			Horizon		
	1-Year	2-Year	5-Year	2-Year	5-Year	10-Year
Avg $\|\Delta w\|$	0.0391	0.0464	0.0515	0.0359	0.0495	0.0889
Δ Sharpe ratio	0.0064 (1.00%)	0.0109 (1.33%)	0.0424 (3.58%)	0.0107 (1.33%)	0.0352 (3.05%)	0.1245 (8.38%)
Equalization fee	0.0011 (3.23%)	0.0028 (4.98%)	0.0144 (10.68%)	0.0022 (4.00%)	0.0123 (9.34%)	0.1241 (33.84%)
	Semi-Annual Rebalancing			Annual Rebalancing		
	Horizon			Horizon		
	2-Year	5-Year	10-Year	5-Year	10-Year	20-Year
Avg $\|\Delta w\|$	0.0304	0.0476	0.089	0.0430	0.0872	0.1161
Δ Sharpe ratio	0.0068 (0.86%)	0.0282 (2.51%)	0.1102 (7.72%)	0.0181 (1.68%)	0.0835 (6.11%)	0.2094 (14.76%)
Equalization fee	0.0014 (2.62%)	0.0100 (7.76%)	0.113 (31.61%)	0.0065 (5.23%)	0.0891 (25.67%)	0.2305 (12.10%)

and Panel B presents the results for conditional policies in which the stock and bond returns each period are scaled by the state variable. Each panel reports the average absolute difference between the approximate versus exact portfolio weights, the difference (as level and percentage in parentheses) in the Sharpe ratios of the two portfolios, and the equalization fee (as level and percentage in parentheses), defined as the yearly fee the investor would be willing to pay to be able to use the exact instead of approximate allocations.

Reviewing the results across horizons, rebalancing frequencies, and panels, it is clear that, consistent with our intuition, the compounding terms are relatively unimportant for short horizons but become important for long horizons. For example, with a 2-year horizon, the equalization fee ranges from four basis points (0.9% of the certainty equivalent) with semiannual rebalancing and no conditioning information to 28 basis points (5% of the certainty equivalent) with monthly rebalancing and conditioning information. Even with a 5-year horizon, the largest equalization fee is 1.4% (10.7% of the certainty equivalent). For horizons beyond 5 years, however, the quality of our approximation deteriorates substantially. With a 10-year horizon, quarterly rebalancing, and conditioning information, for instance, the equalization fee is 12.4%, which constitutes more than one-third of the certainty equivalent of the exact solution.

Analyzing the results more closely reveals some intuitive patterns. The importance of the compounding terms increases with the horizon (holding constant the rebalancing frequency) as well as with the rebalancing frequency (holding constant the horizon). The compounding terms are more important for the conditional polices because these are associated with a higher expected growth rate of wealth and a larger number of compounding terms due to the inclusion of the scaled returns.

We conclude from this experiment that our approach of solving the multiperiod portfolio problem with timing portfolios, which ignore the compounding of excess returns over time, results in little economic loss for short-horizon problems (e.g., a market timing mutual fund with 1- to 5-year horizons) but is far less suitable for long-horizon problems (e.g., a pension fund with 20- to 30-year horizons). For the specific data-generating process in equation (33), the approximation error results in an expected utility loss of 10% or less for horizons up to 5 years. An economic loss of this magnitude seems acceptable given the computational gains that arise from the simplicity of our approach, especially when compared to the usual numerical solutions of multiperiod portfolio problems.

II. Optimal Portfolio Weights Implied by a VAR

Our approach can be data intensive for solving portfolio problems with very long horizons. However, this issue can be overcome by using a statistical model for the returns and state variables. For example, consider a problem with a single risky asset and one conditioning variable, and assume that the log (gross) return and log conditioning variable evolve jointly according to the restricted VAR with normally distributed innovations

$$\begin{bmatrix} \ln R_{t+1} \\ \ln z_{t+1} \end{bmatrix} = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} \ln z_t + \epsilon_{t+1}, \tag{35}$$

where $\epsilon_{t+1} \sim \mathrm{N}[0, \Omega]$. We also assume for simplicity that the risk-free rate is constant.

The dynamics of returns in equation (35) imply the following expanded VAR:

$$\ln Y_{t+1} = A + B \ln Y_t + \nu_{t+1}, \tag{36}$$

where $\ln Y_{t+1} = [\ln R_{t+1}, \ln z_{t+1}, \ln z_t, \ln z_t + \ln R_{t+1}]^\top$ and $\eta_{t+1} \sim \mathrm{N}[0, \Gamma]$ with

$$A = \begin{bmatrix} a_1 \\ a_2 \\ 0 \\ a_1 \end{bmatrix}, B = \begin{bmatrix} 0 & b_1 & 0 & 0 \\ 0 & b_2 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & b_1 + 1 & 0 & 0 \end{bmatrix}, \quad \text{and} \quad \Gamma = \begin{bmatrix} \omega_{11} & \omega_{12} & 0 & \omega_{11} \\ \omega_{12} & \omega_{22} & 0 & \omega_{12} \\ 0 & 0 & 0 & 0 \\ \omega_{11} & \omega_{12} & 0 & \omega_{11} \end{bmatrix}. \tag{37}$$

In Γ, ω_{ij} are the elements of the covariance matrix Ω. The first two unconditional moments of this expanded VAR are given by

$$\begin{aligned} \mu &\equiv \mathrm{E}[\ln Y_{t+1}] = (I - B)^{-1} A \\ \mathrm{vec}(\Sigma) &\equiv \mathrm{vec}(\mathrm{Var}[\ln Y_{t+1}]) = (I - B \otimes B)\mathrm{vec}(\Gamma). \end{aligned} \tag{38}$$

We use this expanded VAR to solve for the moments of returns involved in our solution to the dynamic portfolio choice problem.

A. *Single-Period Problem*

First, consider the single-period portfolio problem. Following equation (12), we construct excess returns on the managed portfolios, that is,

$$\tilde{r}_{t+1} = \left[R_{t+1} - R^f,\, z_t\left(R_{t+1} - R^f\right)\right]^\top. \tag{39}$$

From the extended VAR (38), these returns can be written as

$$\tilde{r}_{t+1} = \Lambda Y_{t+1} + \lambda, \tag{40}$$

where

$$\Lambda = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -R^f & 1 \end{bmatrix} \quad \text{and} \quad \lambda = \begin{bmatrix} -R^f \\ 0 \end{bmatrix}. \tag{41}$$

The optimal single-period portfolio choice for the expanded asset space in equation (16) depends on the first two moments of these returns, which are given by

$$\mathrm{E}[\tilde{r}_{t+1}] = \Lambda \mathrm{E}[Y_{t+1}] + \lambda \quad \text{and} \quad \mathrm{Var}[\tilde{r}_{t+1}] = \Lambda \mathrm{Var}[Y_{t+1}] \Lambda^\top, \tag{42}$$

where, from the joint log-normality of Y_{t+1} and the unconditional moments of the VAR,

$$\begin{aligned} \mathrm{E}\,[Y_{t+1}] &= \exp\left\{\mathrm{E}[\ln Y_{t+1}] + \frac{1}{2}\mathrm{Diag}[\mathrm{Var}[\ln Y_{t+1}]]\right\} \\ \mathrm{Var}[Y_{t+1}] &= (\exp\{\mathrm{Var}[\ln Y_{t+1}]\} - 1)\mathrm{E}[Y_{t+1}]\mathrm{E}[Y_{t+1}]^{\top}. \end{aligned} \tag{43}$$

The moments in equation (42), and hence the optimal portfolio weights, can therefore be evaluated using the unconditional moments of the VAR in equation (38).

B. Multiperiod Portfolio Choice

Next, consider a two-period dynamic problem. The excess returns of the conditional and timing portfolios are

$$\begin{aligned} &\tilde{r}_{t\to t+2} \\ &= [\underbrace{(R_{t+1} - R^f)R^f, z_t(R_{t+1} - R^f)R^f}_{\substack{\text{stocks in period 1,} \\ \text{conditional on } z}}, \underbrace{R^f(R_{t+2} - R^f), R^f z_{t+1}(R_{t+2} - R^f)}_{\substack{\text{stocks in period 2,} \\ \text{conditional on } z}}] \\ &= R^f\left(\begin{bmatrix} \Lambda & 0 \\ 0 & \Lambda \end{bmatrix}\begin{bmatrix} Y_{t+1} \\ Y_{t+2} \end{bmatrix} + \begin{bmatrix} \lambda \\ \lambda \end{bmatrix}\right)^{\top}. \end{aligned} \tag{44}$$

The corresponding first and second moments are

$$\begin{aligned} \mathrm{E}[\tilde{r}_{t\to t+2}] &= R^f\left(\begin{bmatrix} \Lambda & 0 \\ 0 & \Lambda \end{bmatrix}\mathrm{E}\begin{bmatrix} Y_{t+1} \\ Y_{t+2} \end{bmatrix} + \begin{bmatrix} \lambda \\ \lambda \end{bmatrix}\right) \\ \mathrm{Var}[\tilde{r}_{t\to t+2}] &= (R^f)^2\begin{bmatrix} \Lambda & 0 \\ 0 & \Lambda \end{bmatrix}\mathrm{Var}\begin{bmatrix} Y_{t+1} \\ Y_{t+2} \end{bmatrix}\begin{bmatrix} \Lambda^{\top} & 0 \\ 0 & \Lambda^{\top} \end{bmatrix}, \end{aligned} \tag{45}$$

where

$$\begin{aligned} \mathrm{E}\begin{bmatrix} Y_{t+1} \\ Y_{t+2} \end{bmatrix} &= \exp\left\{\mathrm{E}\begin{bmatrix} \ln Y_{t+1} \\ \ln Y_{t+2} \end{bmatrix} + \frac{1}{2}\mathrm{Diag}\left[\mathrm{Var}\begin{bmatrix} \ln Y_{t+1} \\ \ln Y_{t+2} \end{bmatrix}\right]\right\} \\ \mathrm{Var}\begin{bmatrix} Y_{t+1} \\ Y_{t+2} \end{bmatrix} &= \left(\exp\left\{\mathrm{Var}\begin{bmatrix} \ln Y_{t+1} \\ \ln Y_{t+2} \end{bmatrix}\right\} - 1\right)\mathrm{E}\begin{bmatrix} Y_{t+1} \\ Y_{t+2} \end{bmatrix}\mathrm{E}\begin{bmatrix} Y_{t+1} \\ Y_{t+2} \end{bmatrix}^{\top}, \end{aligned} \tag{46}$$

and from the unconditional moments of the VAR

$$\begin{aligned} \mathrm{E}\begin{bmatrix} \ln Y_{t+1} \\ \ln Y_{t+2} \end{bmatrix} &= \begin{bmatrix} \mu \\ \mu \end{bmatrix} \\ \mathrm{Var}\begin{bmatrix} \ln Y_{t+1} \\ \ln Y_{t+2} \end{bmatrix} &= \begin{bmatrix} \Sigma & B\Sigma \\ B\Sigma & \Sigma \end{bmatrix}. \end{aligned} \tag{47}$$

Finally, consider an N-period dynamic problem. Using basic matrix algebra, the excess returns on the conditional and timing portfolios can be written as

$$\tilde{r}_{t\to t+N} = (R^f)^{N-1}\left((I_N \otimes \Lambda)\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix} + (\iota_N \otimes \lambda)\right), \tag{48}$$

where I_N and ι_N denote an N-dimensional identity matrix and vector of ones, respectively. The corresponding first and second moments are

$$\begin{aligned} \mathrm{E}[\tilde{r}_{t\to t+N}] &= (R^f)^{N-1}\left((I_N \otimes \Lambda)\mathrm{E}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix} + (\iota_N \otimes \lambda)\right) \\ \mathrm{Var}[\tilde{r}_{t\to t+N}] &= (R^f)^{2(N-1)}(I_N \otimes \Lambda)\mathrm{Var}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix}\left(I_N \otimes \Lambda^\top\right), \end{aligned} \tag{49}$$

where

$$\begin{aligned} \mathrm{E}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix} &= \exp\left\{ \mathrm{E}\begin{bmatrix} \ln Y_{t+1} \\ \cdots \\ \ln Y_{t+Nfs} \end{bmatrix} + \frac{1}{2}\mathrm{Diag}\left[\mathrm{Var}\begin{bmatrix} \ln Y_{t+1} \\ \cdots \\ \ln Y_{t+N} \end{bmatrix}\right]\right\} \\ \mathrm{Var}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix} &= \left(\exp\left\{\mathrm{Var}\begin{bmatrix} \ln Y_{t+1} \\ \cdots \\ \ln Y_{t+N} \end{bmatrix}\right\} - 1\right)\mathrm{E}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix}\mathrm{E}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix}^\top \end{aligned} \tag{50}$$

and

$$\mathrm{E}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix} = \iota_N \otimes \mu$$

$$\mathrm{Var}\begin{bmatrix} Y_{t+1} \\ \cdots \\ Y_{t+N} \end{bmatrix} = \begin{bmatrix} B^0 & B^1 & B^2 & \cdots & B^{N-1} \\ B^1 & B^0 & B^1 & \cdots & B^{N-2} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ B^{N-1} & B^{N-2} & B^{N-3} & \cdots & B^0 \end{bmatrix} \otimes \Sigma. \quad (51)$$

To summarize, the optimal portfolio weights for the N-period dynamic problem with conditional and timing portfolios, which depend on the first and second moments of the managed portfolio returns, can be evaluated analytically using the coefficient matrix B and the unconditional moments μ and Σ of the VAR (which in turn depend on A, B, and Γ). Since we can estimate the VAR with a relatively modest time series of returns and state variable realizations, we can solve dynamic portfolio choice problems with arbitrarily long horizons using finite data samples in this VAR context. Of course, this comes at the cost of having to impose strong structure on the dynamics of returns.

III. Extensions and Refinements

One can extend and refine our approach along a number of dimensions. In this section, we show how to generalize the investor's utility function and how to compute robust portfolio weights for a large number of assets using techniques developed originally for the static Markowitz approach.

A. Objective Functions

The mean-variance objective function can be extended to an arbitrary utility function $u(W_{t+1})$. In that case, we solve the problem

$$\max_{\theta} \mathrm{E}_t\big[u\big(R_t^f + (\theta z_t)^\top r_{t+1}\big)\big], \quad (52)$$

or the corresponding first-order conditions, using numerical optimization methods. While high-dimensional numeric solutions are nontrivial, our approach benefits from being static and unconstrained (since we ignore the compounding terms). Furthermore, the extensive literature on effective and fast algorithms for solving high-dimensional optimization problems applies to our framework. These algorithms include variants of the Newton method (e.g., Conn, Gould, and Toint (1988), Moré and Toraldo (1989)), the quasi-Newton method (e.g., Byrd et al. (1995)), and the sequential quadratic programming approach (e.g., Gill, Murray, and Saunders (2002)).

The quadratic objective function (4) can be interpreted alternatively as a second-order approximation of a more general utility function, such as power or more general hyperbolic absolute risk aversion (HARA) preferences. To increase the precision of this approximation, Brandt et al. (2005) propose a fourth-order expansion that includes adjustments for the skewness and kurtosis of returns, and their effects on expected utility. Specifically, the expansion of expected utility around the current wealth growing at the risk-free rate is

$$\mathrm{E}_t\left[u(W_{t+1})\right] \approx \mathrm{E}_t\left[u\left(W_t R_t^f\right) + u'\left(W_t R_t^f\right)\left(W_t x_t^\top r_{t+1}\right) + \frac{1}{2}u''\left(W_t R_t^f\right)\left(W_t x_t^\top r_{t+1}\right)^2 + \frac{1}{6}u'''\left(W_t R_t^f\right)\left(W_t x_t^\top r_{t+1}\right)^3 + \frac{1}{24}u''''\left(W_t R_t^f\right)\left(W_t x_t^\top r_{t+1}\right)^4\right]. \quad (53)$$

In this case, the first-order conditions define an implicit solution for the optimal weights in terms of the joint moments of the derivatives of the utility function and returns, given by,

$$x_t \approx -\left\{\mathrm{E}_t\left[u''\left(W_t R_t^f\right)\left(r_{t+1} r_{t+1}^\top\right)\right] W_t^2\right\}^{-1} \times \left\{\mathrm{E}_t\left[u'\left(W_t R_t^f\right)(r_{t+1})\right] W_t + \frac{1}{2}\mathrm{E}_t\left[u'''\left(W_t R_t^f\right)\left(x_t^\top r_{t+1}\right)^2 r_{t+1}\right] W_t^3 + \frac{1}{6}\mathrm{E}_t\left[u''''\left(W_t R_t^f\right)\left(x_t^\top r_{t+1}\right)^3 r_{t+1}\right] W_t^4\right\}. \quad (54)$$

In practice, this implicit expression for the optimal weights is easy to solve. Start with an initial "guess" for the optimal weights (such as equal weights in each asset), denoted by $x_t(0)$. Then enter this guess on the right-hand side of equation (54) and obtain a new solution for the optimal weights on the left-hand side, denoted by $x_t(1)$. After a few iterations n, the guess $x_t(n)$ is very close to the solution $x_t(n+1)$ and we can take this value to be the solution of equation (54). Brandt et al. (2005) show that this expansion is highly accurate for investment horizons up to 1 year, even when returns are far from normally distributed. Use of this expansion approach in our extended asset space approach is straightforward.

We can also consider performance benchmarks in the objective function. Frequently, money managers are evaluated on their performance relative to a benchmark index portfolio over a given period. Such problems can be solved easily with our approach. Simply use returns of the basis assets in excess of the benchmark index (instead of in excess of the risk-free rate) in the portfolio optimization. In this case, the objective function involving these excess returns captures the gain from beating the benchmark index with low tracking error. The optimal portfolio weights can be interpreted as deviations from the benchmark, which are usually termed as "active" weights.

Finally, we can expand the mean-variance objective to penalize covariance with the return of a particular portfolio such as the market index, denoted by r^m.[4] In this case, the objective is

[4] Or, similarly, penalize covariance with consumption growth.

$$\mathrm{E}_t\left[r^p_{t+1} - \frac{\gamma}{2}\left(r^p_{t+1}\right)^2 - \lambda r^p_{t+1} r^m_{t+1}\right], \tag{55}$$

with some positive penalty constant λ. Replacing population moments by sample moments, the solution in the unconditional case is

$$x = \frac{1}{\gamma}\left[\sum_{t=1}^{T-1} r_{t+1} r_{t+1}^{\top}\right]^{-1}\left[\sum_{t=1}^{T-1}\left(1 - \lambda r^m_{t+1}\right) r_{t+1}\right], \tag{56}$$

which can be extended trivially to the conditional and multiperiod problems.

B. Constraints, Shrinkage, and Prior Views

A benefit of framing the dynamic portfolio problem in a static context is that we have available all of the refinements of the Markowitz approach that have been developed over the past decades. These include the use of portfolio constraints to avoid extreme positions (e.g., Frost and Savarino (1988), Jagannathan and Ma (2003)), the use of shrinkage to improve the estimates of the means (e.g., Jobson and Korkie (1981)) as well as of the covariance matrix (e.g., Ledoit (1995)), and the combination of the investor's prior from an alternative data source or the belief in a pricing model with the information contained in returns (e.g., Treynor and Black (1973), Black and Litterman (1992), and Pastor and Stambaugh (2000)).

For the last approach, which is particularly useful in practice, a natural prior is that the market is in equilibrium. In that case the market portfolio is the tangency portfolio. Suppose that the estimated portfolio weight on asset i is of the form $x^i_t = a + bz_t$, and assume that z has been standardized to have mean zero. Using the equilibrium prior, we would shrink a toward the market capitalization weight of the asset and b toward zero. The shrinkage weights can be determined from the standard errors of the estimates of a and b, coupled with a prior on the efficiency of the market.

IV. Application

There is substantial evidence that economic variables related to the business cycle help forecast stock and bond returns. For instance, Campbell (1991), Campbell and Shiller (1988), Fama (1990), Fama and French (1988, 1989), Hodrick (1992), and Keim and Stambaugh (1986) report evidence that stock market returns are predicted by the dividend-price ratio, short-term interest rate, term spread, and credit spread. Fama and French (1989) show that the same variables also predict bond returns. We use these four conditioning variables in a simple application of our method to the dynamic portfolio choice among stocks, bonds, and cash. This application is similar to that of Brennan, Schwartz, and Lagnado (1997) and Campbell, Chan, and Viceira (2003).

We use as proxy for stocks the CRSP value-weighted market index, for bonds the index of long-term Treasuries constructed by Ibbotson Associates, and for cash the 1-month Treasury bill, also obtained from Ibbotson Associates. The dividend-price ratio (*DP*) is calculated as the difference between the log of the last 12 months of dividends and the log of the current price of the CRSP value-weighted index. The relative Treasury bill (*Tbill*) stochastically detrends the raw series by taking the difference between the Treasury bill rate and its 12-month moving average. The term spread (*Term*) is the difference between the yields on 10-year and 1-year government bonds. The default spread (*Default*) is calculated as the difference between the yield on BAA- and AAA-rated corporate bonds. We obtain interest rate data from the DRI/Citibase database, and we standardize the four predictors to ease the interpretation of the coefficients of the portfolio policy. The sample period is January 1945 through December 2000.

Table II reports the results for both unconditional and conditional portfolio policies at monthly, quarterly, and annual holding periods, assuming the investor has quadratic utility with $\gamma = 5$. Some differences in the unconditional portfolio weights exist across the three holding periods. With monthly or quarterly rebalancing, the weight in equities is 77%, whereas it is only 57% with annual rebalancing. This pattern is due to differences in the joint distribution of stock and bond excess returns over the different holding periods. In particular, the small amount of positive serial correlation in returns at the monthly and quarterly frequencies turns negative at the annual frequency, which makes the volatility of stock and bond returns proportionately higher at the annual frequency (15.6% vs. 14.5% for stocks and 9.8% vs. 8.4% for bonds). The weight invested in bonds is close to zero for all holding periods, so the investor allocates roughly 25% to 45% to the risk-free asset.

The conditional policies are quite sensitive to the state variables. For the monthly conditional policy, the coefficients of the stock weight on *Default* and *DP* as well as the coefficients of the bond weight on *DP* and *TBill* are all significant at the 5% level. Furthermore, the average allocations to stocks and bonds by the conditional policy are 87% and 29%, respectively, which significantly exceed the corresponding unconditional allocations of 77% and −1%. The reason is that the predictability in the first and second moments of returns allows the investor to be more aggressive on average since the exposure can be reduced in bad times (i.e., times in which the mean return of the optimal portfolio is low and/or its volatility is high). An F-test of the hypothesis that all coefficients on the state variables are equal to zero has a p-value of zero. Finally, the (annualized) Sharpe ratio of the conditional policy is 1.00, which is nearly 70% higher than that of the unconditional policy of 0.59. Overall, it is clear that on average the conditional return distribution is very different from the unconditional return distribution.

The results are less pronounced for the longer holding periods. At the quarterly horizon, for example, only the coefficients of the bond weight on *Term* are significant at the 5% level (the coefficients of the stock weight on *Term* and *TBill*

Table II

Single-Period Portfolio Policies

This table provides estimates of the single-period portfolio policy. Standard errors for the coefficients of the portfolio policies are in parentheses. The p-value refers to an F-test of the hypothesis that all the coefficients on the state variables other than the constant are jointly zero. The next three rows present statistics of the returns generated by the portfolio policies. The data are from January of 1945 to December of 2000. Finally, the last row shows the equalization fee, defined as the yearly fee that the investor would be willing to pay to be able to use the conditional instead of the unconditional policy.

	State Variable	Monthly		Quarterly		Annual	
		Unconditional	Conditional	Unconditional	Conditional	Unconditional	Conditional
Stock	Const	0.764 (0.185)	0.873 (0.186)	0.770 (0.179)	0.649 (0.185)	0.572 (0.139)	0.591 (0.177)
	Term		−0.166 (0.184)		0.304 (0.223)		0.249 (0.203)
	Default		−0.504 (0.173)		−0.333 (0.183)		0.046 (0.155)
	D/P		0.796 (0.306)		0.126 (0.190)		0.037 (0.149)
	TBill		−0.485 (0.459)		−0.434 (0.235)		−0.172 (0.242)
Bond	Const	−0.005 (0.323)	0.291 (0.174)	−0.040 (0.319)	−0.084 (0.389)	−0.074 (0.255)	−0.696 (0.323)
	Term		−0.096 (0.173)		0.840 (0.336)		0.352 (0.278)
	Default		−0.465 (0.404)		0.144 (0.351)		0.337 (0.291)
	D/P		0.672 (0.310)		−0.076 (0.458)		−0.264 (0.373)
	TBill		0.773 (0.247)		0.601 (0.346)		−0.152 (0.397)
p-value			0.000		0.000		0.000
Mean excess return		0.063	0.185	0.065	0.126	0.049	0.095
SD return		0.111	0.185	0.109	0.146	0.087	0.101
Sharpe ratio		0.589	1.000	0.602	0.864	0.571	0.938
Equalization fee		0.067		0.037		0.039	

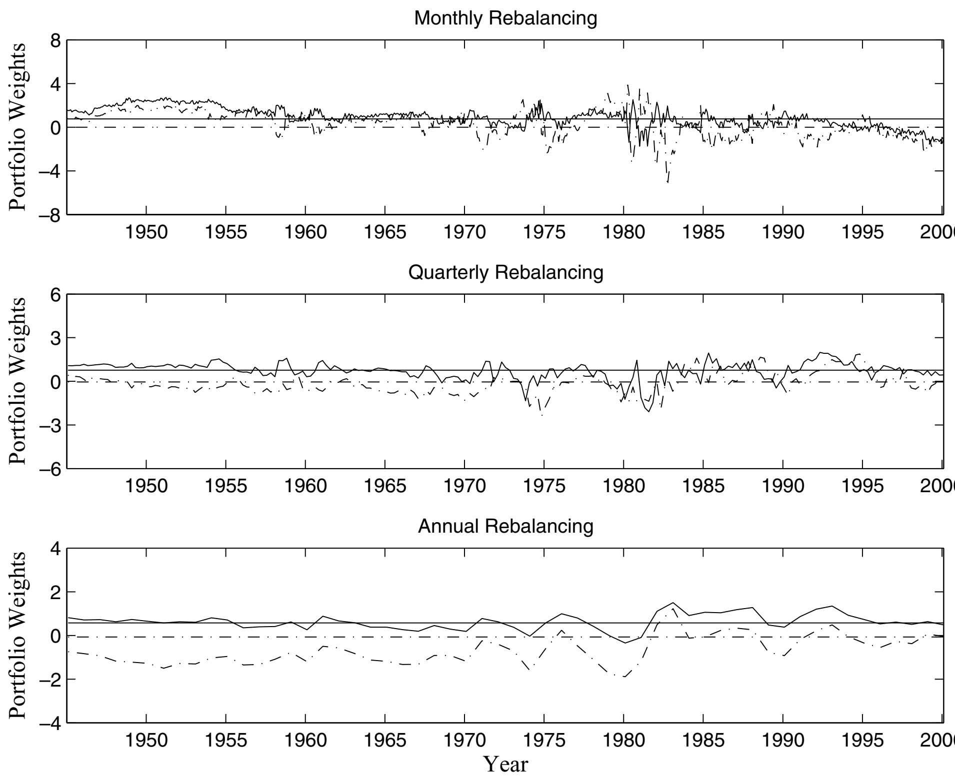

Figure 1. Portfolio weights of conditional and unconditional policies. This figure displays the time series of conditional portfolio weights. The solid line corresponds to the portfolio weight on the stock and the dash–dotted line corresponds to the portfolio weight on the bond. The constant portfolio weights from the unconditional policy are depicted as straight lines.

are significant at the 10% level). However, the hypothesis that all coefficients are zero is still rejected with a p-value of zero. More importantly, the Sharpe ratio of the conditional policy is still 40% higher than that of the unconditional policy, with a value of 0.86 versus 0.60. The results for the annual policy are qualitatively similar, with an increase in the Sharpe ratio from 0.57 to 0.94 due to conditioning.

Figure 1 displays the time series of portfolio weights of the conditional policies. For comparison, the figure also shows the unconditional portfolio weights. Overall, the shorter the holding period, the more extreme positions the policies take at times (note the different scales on the y-axis). It is striking that the conditional policies can be substantially different at different frequencies.

The most striking difference in the portfolio policies across horizons lies in the average bond holdings (corresponding to the bond intercepts of the portfolio policy). With monthly rebalancing, the optimal conditional allocation to bonds is 29%. With annual rebalancing, in contrast, the portfolio is short 70% bonds.

These extreme differences in portfolio weights are due to drastic changes in the conditional volatilities and correlations of bond and stock returns across horizons. At a monthly frequency, the average annualized conditional volatilities of stock and bond returns are 11.2% and 5.8%, respectively, with an average conditional correlation of 0.22. At the annual frequency, the average conditional volatility of stock returns increases to 12.5% while the average conditional volatility of bond returns drops to 5.5%. Most importantly, the average conditional correlation increases to 0.37. Given the positive correlation, a risk-averse investor shorts bonds at long horizons to diversify. And since the risk premium on bonds is very small relative to that of stocks (1% compared to more than 8% in our sample), shorting bonds is not very costly. This effect is stronger at the annual frequency given the higher average conditional correlation. In addition to this diversification effect, with annual rebalancing, the portfolio policy is not sufficiently responsive to the predictors to sell bonds during the short periods of time in the early 1980s when bond returns were extremely volatile and negative. With monthly rebalancing, in contrast, the portfolio policy is flexible enough to be long bonds at the beginning and end of the sample, while being very short bonds during periods in the early 1980s. In some sense, in order to be short bonds in the 1980s, the annual policy needs to also be short bonds in other time periods, making the bond holdings negative on average. The monthly policy is more flexible. It is able to take negative bond positions in bad months while on average holding a positive weight in bonds. This difference between the monthly and yearly bond position is apparent in the volatility of the bond portfolio weights exhibited in the top plot of Figure 1.

By focusing directly on the portfolio weights we capture time variation in the entire return distribution as opposed to just the expected returns. To get a sense of the importance of this aspect of our approach, we compare the conditional policies to more traditional strategies based only on predictive return regressions. Specifically, we regress the excess stock and bond returns on the state variables and then use the corresponding one-period-ahead forecasts of the returns together with the unconditional covariance matrix to form portfolio weights. In this way, the strategy only takes into account the predictability of expected returns and ignores the impact of the state variables on variances and covariances. Table III compares the two approaches, and Figure 2 plots the time series of portfolio weights on the stock.

The advantage of our approach is most apparent at the monthly frequency. Although our conditional strategy generates a lower premium of 18.5% versus 25.5% per year, it has proportionally much lower volatility of 18.5% versus 28.6% per year, resulting in a Sharpe ratio that is 12% higher (1.01 compared to 0.89). In fact, the investor is willing to pay an annual fee of 5.6% to obtain the improved performance associated with exploiting the joint time variation of the entire return distribution, as opposed to using the time variation of the mean returns only.

To get a more clear sense of where this performance improvement is coming from, compare the coefficients of the monthly portfolio policy in Table II to the regression coefficients in Table III. The most striking difference is that *Default*

Table III

Traditional versus Optimal Conditional Policies

This table reports estimates of the traditional approach to tactical asset allocation. In this approach, conditional expected returns are obtained from an in-sample regression of returns on the state variables and the Markowitz solution is applied to these conditional expected returns together with the unconditional covariance matrix. Panel A displays the estimated regressions of stock and bond returns on the conditioning variables, each estimated at the monthly, quarterly, and annual frequency. Panel B summarizes the traditional portfolio policy (Trdnl) and, for comparison, the full conditional policy (Cndtnl) that takes into account the impact of the conditioning variables both on expected returns and their covariance matrix. The first two rows present the time-series average of the weights on stocks and bonds of the two policies. The next three lines offer statistics for the time series of portfolio returns. The last row shows the equalization fee, defined as the yearly fee that the investor would be willing to pay to be able to use the full conditional policy instead of using the traditional approach.

Panel A: Regression Estimates

	Coefficient	Monthly	Quarterly	Annual
Stock	Cnst	0.0829 (0.0191)	0.0779 (0.0188)	0.0838 (0.0213)
	Term	0.0276 (0.0230)	0.0340 (0.0224)	0.0336 (0.0244)
	Default	0.0116 (0.0204)	0.0061 (0.0207)	−0.0061 (0.0229)
	DP	−0.0206 (0.0195)	0.0144 (0.0194)	0.0340 (0.0219)
	Tbill	−0.0756 (0.0238)	−0.0534 (0.0259)	−0.0217 (0.0335)
	R^2	0.0384	0.0718	0.1139
Bond	Cnst	0.0120 (0.0111)	0.0134 (0.0109)	0.0138 (0.0113)
	Term	0.0586 (0.0133)	0.0524 (0.0129)	0.0553 (0.0130)
	Default	0.0424 (0.0118)	0.0275 (0.0120)	0.0178 (0.0122)
	DP	−0.0087 (0.0113)	−0.0017 (0.0112)	−0.0025 (0.0116)
	Tbill	0.0373 (0.0138)	0.0233 (0.0149)	0.0023 (0.0177)
	R^2	0.0432	0.0919	0.365

Panel B: Portfolio Policies

	Monthly		Quarterly		Annual	
	Trdnl	Cndtnl	Trdnl	Cndtnl	Trdnl	Cndtnl
Mean weight stock	0.7833	0.8728	0.8342	0.6486	0.7484	0.5914
Mean weight bond	−0.0052	0.2911	−0.0436	−0.0842	−0.0969	−0.6958
Mean excess return	0.2552	0.1849	0.1877	0.1262	0.1489	0.0945
SD return	0.2868	0.1849	0.2617	0.1461	0.1909	0.1008
Sharpe ratio	0.8901	1.0001	0.7172	0.8638	0.7801	0.9375
Equalization fee	0.0560		0.0679		0.0432	

and *DP* have opposite signs in the regressions (for both stocks and bonds, the *Default* coefficient is positive and the *DP* coefficient is negative) compared to the conditional portfolio weights. Furthermore, only *Default* is significant in the regression for bonds, while *Default* and *DP* are significant (with opposite signs) in the stock portfolio weight and *DP* is significant (again, with opposite sign) in the bond portfolio weight. The reason is that *Default* and *DP* are significant predictors of volatility, in particular of bond return volatility. Together, *Default* and *DP* explain 2.3% of absolute stock returns and 21% of absolute bond returns

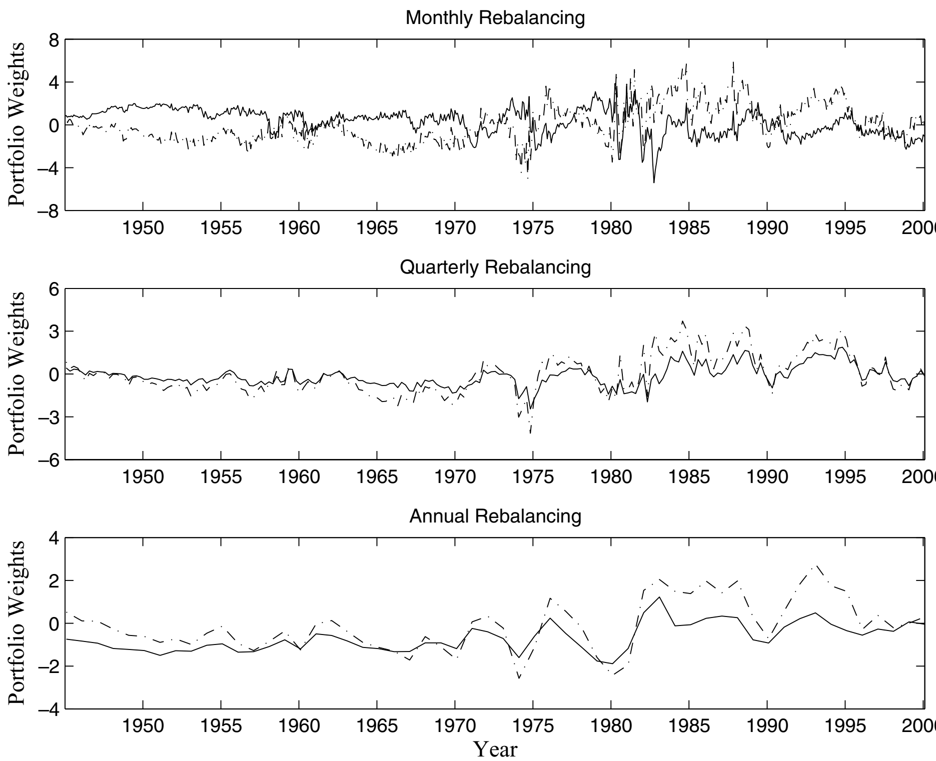

Figure 2. Portfolio weights of conditional and regression-based policies. This figure displays the time series of the portfolio weight on the stock obtained from the conditional approach (solid line) and from the regression-based approach (dashed line). In the regression-based approach, conditional-expected returns are computed from an in-sample regression of returns on the state variables, and the Markowitz solution is applied to these conditional-expected returns together with the unconditional covariance matrix.

in our sample, with positive coefficients on *Default* and negative coefficients on *DP*. Consequently, the conditional portfolio weights load strongly negatively on *Default* and strongly positively on *DP*, resulting in very different allocations relative to the traditional policy, in particular for bonds (average holding of 29% as opposed to −1%).

Although the differences between the two strategies are less dramatic at lower frequencies, the conclusion holds nevertheless. The fee the investor is willing to pay for using the conditional strategy as opposed to the regression approach is 6.8% with quarterly rebalancing and 4.3% with annual rebalancing. Furthermore, the differences in the signs of the coefficients are all explained by the predictive power of the state variables for the second moments of stock and bond returns.

We now turn our attention to multiperiod strategies. Table IV reports the portfolio weights of the multiperiod portfolio policy for a 1-year horizon with monthly or quarterly rebalancing. For simplicity, we report only the

Table IV

Multiperiod Portfolio Policies

This table shows estimates of the multiperiod portfolio policy with a 1-year horizon and monthly or quarterly rebalancing. Standard errors for the coefficients of the portfolio policies are in parentheses. The p-value refers to an F-test of the hypothesis that all the coefficients on the state variables other than the constant are jointly zero. The next three rows present statistics of the returns generated by the portfolio policies. Finally, the last row shows the equalization fee, defined as the yearly fee that the investor would be willing to pay to be able to use the conditional instead of the unconditional policy.

			Monthly		Quarterly	
Asset	Month/ Quarter	State Variable	Unconditional	Conditional	Unconditional	Conditional
Stock	1/1	Cnst	0.6276 (0.1602)	0.6779 (0.1598)	0.6332 (0.1188)	0.6200 (0.1622)
		Tbill		−0.6908 (0.2834)		−0.2541 (0.1499)
	4/2	Cnst	0.6205 (0.1645)	0.5810 (0.1621)	0.6169 (0.1171)	0.5750 (0.1656)
		Tbill		−0.2905 (0.2856)		−0.2059 (0.1565)
	8/3	Cnst	0.5508 (0.1671)	0.5191 (0.1653)	0.5505 (0.1172)	0.5587 (0.1688)
		Tbill		0.0794 (0.2814)		−0.2569 (0.1622)
	12/4	Cnst	0.4837 (0.1645)	0.4138 (0.1617)	0.5461 (0.1198)	0.4104 (0.1649)
		Tbill		0.3763 (0.2823)		−0.3857 (0.1577)
Bond	1/1	Cnst	−0.6860 (0.2891)	−0.2978 (0.1390)	−0.5543 (0.2901)	−0.7115 (0.2971)
		Tbill		0.0805 (0.1729)		0.2125 (0.2200)
	4/2	Cnst	−0.3405 (0.2906)	−0.1860 (0.1432)	−0.2712 (0.2942)	−0.2843 (0.2985)
		Tbill		−0.0424 (0.1731)		0.1940 (0.2256)
	8/3	Cnst	0.0329 (0.2875)	−0.1290 (0.1438)	−0.0256 (0.1873)	−0.0623 (0.3010)
		Tbill		−0.0968 (0.1735)		−0.0078 (0.2103)
	12/4	Cnst	0.4628 (0.2849)	−0.4408 (0.1409)	0.2923 (0.2855)	0.2884 (0.2989)
		Tbill		0.0161 (0.1715)		0.0151 (0.2093)
p-value			0.0000	0.0000	0.0000	0.0000
Mean excess return			0.0507	0.0687	0.0526	0.0658
SD return			0.0871	0.0951	0.0883	0.0942
Sharpe ratio			0.5824	0.7224	0.4740	0.6985
Equalization fee			0.0143		0.0105	

unconditional strategy and the conditional strategy with a single state variable, the detrended T-bill rate. The table reports the estimated portfolio weights for month 1, 4, 8, and 12 as well as for all four quarters of the 12-month or four-quarter problems.

With monthly rebalancing, the weight on stocks decreases and the weight on bonds increases as the end of the horizon approaches. This horizon pattern is roughly the same for the unconditional and conditional policies, which means that it is generated by the serial-covariance structure of the returns on the basis assets. With quarterly rebalancing, the unconditional and average conditional (the constant term in the conditional policy) stock holdings are similar to each other and to the results with monthly rebalancing. The unconditional and average conditional bond holdings, in contrast, are very different from each other. In the unconditional policy, the bond holding increases from −69% to 46% as the end of the horizon approaches, while in the conditional policy the average bond holding decreases from −30% to −44% percent.[5] This difference in the horizon

[5] The multiperiod allocations cannot be directly compared to the single-period ones. In particular, the positive bond holdings in the single-period problem are driven by the bond return predictability from *Default* and *DP*, both of which are not included in the multiperiod problem.

patterns can only be attributed to the serial-covariance structure of the conditional portfolio returns, which illustrates the importance of augmenting the asset space in this multiperiod problem.

V. Conclusion

We present a simple approach for dynamic portfolio selection. The model extends the Markowitz approach to the choice between managed portfolios, specifically, between conditional portfolios that invest in each asset an amount that is proportional to some conditioning variable and timing portfolios that invest in each asset for a single period and in the risk-free asset for all other periods. The intuition underlying our approach is that the static choice among these mechanically managed portfolios is equivalent to a dynamic strategy in the basis assets. Our hope is that, by making dynamic portfolio selection no more difficult to implement than the static Markowitz approach, it will finally leave the confines of the ivory tower and make its way into the day-to-day practice of the investment industry.

REFERENCES

Aït-Sahalia, Yacine, and Michael W. Brandt, 2001, Variable selection for portfolio choice, *Journal of Finance* 56, 1297–1351.

Aït-Sahalia, Yacine, and Michael W. Brandt, 2005, Portfolio and consumption choice with option-implied state prices, Working paper, Princeton University.

Bansal, Ravi, and Campbell R. Harvey, 1996, Performance evaluation in the presence of dynamic trading strategies, Working paper, Duke University.

Black, Fisher, and Robert Litterman, 1992, Global portfolio optimization, *Financial Analysts Journal* 48, 28–43.

Brandt, Michael W., 1999, Estimating portfolio and consumption choice: A conditional method of moments approach, *Journal of Finance* 54, 1609–1646.

Brandt, Michael W., Amit Goyal, Pedro Santa-Clara, and Jonathan R. Stroud, 2005, A simulation approach to dynamic portfolio choice with an application to learning about return predictability, *Review of Financial Studies* 18, 831–873.

Brandt, Michael W., Pedro Santa-Clara, and Rossen Valkanov, 2005, Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns, Working paper, Duke University.

Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.

Britten-Jones, Mark, 1999, The sampling error in estimates of mean-variance efficient portfolio weights, *Journal of Finance* 54, 655–671.

Byrd, Richard H., Peihuang Lu, Jorge Nocedal, and Ciyou Zhu, 1995, A limited memory algorithm for bound constrained optimization, *SIAM Journal on Scientific Computing* 16, 1190–1208.

Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.

Campbell, John Y., Yeung Lewis Chan, and Luis M. Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics* 67, 41–80.

Campbell, John Y., and Robert J. Shiller, 1988, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.

Campbell, John Y., and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics* 114, 433–495.

Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press, New York).

Conn, Andy R., Nick I.M. Gould, and Philippe Toit, 1988, Global convergence for a class of trust-region algorithms for optimization with simple bounds, *SIAM Journal of Numerical Analysis* 25, 433–460.

Cox, John C., and Chi-Fu Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of Economic Theory* 49, 33–83.

Fama, Eugene F., 1990, Stock returns, expected returns, and real activity, *Journal of Finance* 45, 1089–1108.

Fama, Eugene F., and Kenneth R. French, 1988, Permanent and temporary components of stock prices, *Journal of Political Economy* 96, 246–273.

Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.

Ferson, Wayne E., and Andrew F. Siegel, 2001, The use of conditioning information in portfolios, *Journal of Finance* 56, 967–982.

Frost, Peter A., and James E. Savarino, 1988, For better performance: Constrain portfolio weights, *Journal of Portfolio Management* 15, 29–34.

Gill, Philip E., Walter Murray, and Michael A. Saunders, 2002, SNOPT: An SQP algorithm for large-scale constrained optimization, *SIAM Journal of Optimization* 12, 979–1006.

Hansen, Lars P., and Scott F. Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.

Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 257–286.

Jagannathan, Ravi, and Tongshu Ma, 2003, Risk reduction in large portfolios: Why imposing the wrong constraints helps, *Journal of Finance* 58, 1651–1683.

Jobson, J. David, and Bob Korkie, 1981, Putting Markowitz theory to work, *Journal of Portfolio Management* 7, 70–74.

Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357–390.

Ledoit, Olivier, 1995, A well-conditioned estimator for large dimensional covariance matrices, Working paper, UCLA.

Moré, Jorge J., and Gerardo Toraldo, 1989, Algorithms for bound constrained quadratic programming problems, *Numerical Mathematics* 55, 377–400.

Pastor, Lubos, and Robert F. Stambaugh, 2000, Comparing asset pricing models: An investment perspective, *Journal of Financial Economics* 56, 335–381.

Treynor, Jack L., and Fisher Black, 1973, How to use security analysis to improve portfolio selection, *Journal of Business* 46, 66–86.