

NBER WORKING PAPER SERIES

OPTIMAL VALUE AND GROWTH TILTS
IN LONG-HORIZON PORTFOLIOS

Jakub W. Jurek
Luis M. Viceira

Working Paper 12017
<http://www.nber.org/papers/w12017>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2006

Jurek: Harvard Business School, Boston MA 02163. Email: jjurek@hbs.edu. Viceira: Harvard Business School, Boston MA 02163, CEPR and NBER. Email: lviceira@hbs.edu. We are grateful for helpful comments and suggestions by John Campbell, Domenico Cuoco, Wayne Ferson, Stavros Panageas, Raman Uppal, Jessica Wachter, Motohiro Yogo, and seminar participants at Harvard Business School, the Wharton School of the University of Pennsylvania, the Kenan-Flagler Business School of the University of North Carolina at Chapel Hill, the Finance Department at Boston College, the Sauder School of Business of the University of British Columbia, and the 2006 North American Winter Meeting of the Econometric Society. We also thank the Division of Research at HBS for generous financial support. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Jakub W. Jurek and Luis M. Viceira. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Value and Growth Tilts in Long-Horizon Portfolios

Jakub W. Jurek and Luis M. Viceira

NBER Working Paper No. 12017

January 2006

JEL No. G12

ABSTRACT

We develop an analytical solution to the dynamic portfolio choice problem of an investor with power utility defined over wealth at a finite horizon who faces an investment opportunity set with time-varying risk premia, real interest rates and inflation. The variation in investment opportunities is captured by a flexible vector autoregressive parameterization, which readily accommodates a large number of assets and state variables. We find that the optimal dynamic portfolio strategy is an affine function of the vector of state variables describing investment opportunities, with coefficients that are a function of the investment horizon. We apply our method to the optimal portfolio choice problem of an investor who can choose between value and growth stock portfolios, and among these equity portfolios plus bills and bonds. For equity-only investors, the optimal mean allocation of short-horizon investors is heavily tilted away from growth stocks regardless of their risk aversion. However, the mean allocation to growth stocks increases dramatically with the investment horizon, implying that growth is less risky than value at long horizons for equity-only investors. By contrast, long-horizon conservative investors who have access to bills and bonds do not hold equities in their portfolio. These investors are concerned with interest rate risk, and empirically growth stocks are not particularly good hedges for bond returns. We also explore the welfare implications of adopting the optimal dynamic rebalancing strategy vis a vis other intuitive, but suboptimal, portfolio choice schemes and find significant welfare gains for all long-horizon investors.

Jakub W. Jurek

Graduate School of Business Administration

Harvard University

Boston, MA 02163

jjurek@hbs.edu

Luis M. Viceira

Graduate School of Business Administration

Harvard University

Boston, MA 02163

and NBER

lviceira@hbs.edu

1 Introduction

Long-term investors seek portfolio strategies that optimally trade off risk and reward, not in the immediate future, but over the long term. Consider for example a long-term investor who cares only about the distribution of her wealth at some given future date. Today, at time t , the investor picks a portfolio to maximize the expected utility of wealth at time $t + K$, where K is the investment horizon. If the portfolio must be chosen once and for all, with no possibility of rebalancing between t and $t + K$, then this is a static portfolio choice problem of the sort studied by Markowitz (1952). The solution depends on the risk properties of returns measured over K periods, but given these risk properties the portfolio choice problem is straightforward.²

It is unrealistic, however, to assume that long-term investors can be expected to adopt this “invest and forget” strategy. Investors typically choose portfolio policies that require periodic rebalancing of portfolio weights. For example, practitioners often advocate the use of a constant proportion strategy in which the portfolio is rebalanced each period to a fixed vector of weights.

More generally investors might optimally choose to rebalance their portfolios to a vector of portfolio weights that adapts to changing market conditions between t and $t + K$ as investment opportunities (i.e. risk premia, interest rates, inflation, etc.) vary over time. In this case the investor must find, not a single optimal portfolio, but an optimal dynamic portfolio strategy or contingent plan that specifies how to adjust asset allocations in response to the changing investment opportunity set. Solving for this contingent plan is a challenging problem. Samuelson (1969) and Merton (1969, 1971, 1973) showed how to use dynamic programming to characterize the solution to this type of problem, but did not derive closed-form solutions except for the special cases where the long-term strategy is identical to a sequence of optimal short-term strategies.

In recent years financial economists have explored many alternative solution methods for the long-term portfolio choice problem with rebalancing. Exact analytical solutions have been discovered for a variety of special cases (e.g. Kim and Omberg 1996, Liu 1998, Brennan and Xia 2002, Wachter 2002), but these often fail to capture all the dimensions of variation in the investment opportunity set that appear to be relevant empirically. In particular, these models typically do not allow both the real interest rate and risk premia to vary over time. Numerical methods have also been developed for this type of problem and range from discretizing the state space (e.g. Balduzzi and Lynch 1999, Barberis 2000) to numerically

²Campbell and Viceira (2005) provide an accessible discussion of the risk properties of US stocks, bonds, and Treasury bills at long horizons and the implications for optimal long-term buy-and-hold portfolios.

solving the PDE characterizing the dynamic program (e.g. Brennan, Schwartz, and Lagnado 1997, 1999). Although numerical methods can, in principle, handle arbitrarily complex model setups with realistic return distributions and portfolio constraints, in practice it has proven difficult to use these methods in problems with more than a few state variables. Finally, there are approximate analytical methods that deliver solutions that are accurate in the neighborhood of special cases for which closed-form results are available. Campbell and Viceira (1999, 2001, 2002) develop this approach for the case of an infinitely lived investor who derives utility from consumption rather than wealth. Their method is accurate provided the investor’s consumption-wealth ratio is not too variable. Campbell, Chan, and Viceira (2003, CCV henceforth) apply the method to a problem with multiple risky assets and allow both the real interest rate and risk premia to change over time.

This paper makes several contributions to the portfolio choice literature. First, we provide an analytical recursive solution to the dynamic portfolio choice problem of an investor whose utility is defined over wealth at a single fixed horizon, in an environment with time-varying investment opportunities. The variation in risk premia, inflation, interest rates and the state variables that drive them is captured using a VAR(1) model. By using the vector autoregressive framework we are able to conveniently handle a large number of investable assets and state variables, which often pose significant problems for numerical methods. Our recursive solution is based on the Campbell-Viceira approximation to the log-portfolio return, and consequently—like the approximation itself—is exact in continuous time. In this sense, it can be interpreted as a generalization of the solutions in Kim and Omberg (1996), Brennan and Xia (2002), and others, to an arbitrarily intricate state-space. This novel solution allows us to examine horizon effects in portfolio choice. In concurrent work, Sorensen and Trolle (2005) derive a solution similar to ours, which they use to study dynamic asset allocation with latent state variables.

Second, we compare this solution with a simpler portfolio rule that would be optimal if rebalancing were restricted. We consider a K -period constant proportion strategy that rebalances the portfolio weights every period to a fixed vector which is optimally chosen by investors at the beginning of their investment horizon, taking into account that time variation in risk premia and interest rates creates a term structure of expected returns and risk. This case provides useful intuition about the link between unconditional geometric and arithmetic mean returns, and horizon effects on portfolio choice. It also serves as the alternative to the dynamically rebalanced strategy in our welfare analysis, where we document that access to the optimal dynamically rebalanced strategy provides a large benefit relative to the constant proportion strategy.

Finally, we apply our method to an empirically relevant problem: optimal growth and value tilts in the portfolios of long-horizon investors. Most studies of empirically motivated

optimal dynamic portfolio choice problem focus on the choice between a well-diversified portfolio of equities representing the market, and other assets such as cash and long-term bonds. These studies artificially constrain investors who want exposure to equities to hold the aggregate stock market portfolio. However, in an environment of changing expected returns, it is empirically plausible that investors optimally choose equity portfolios which do not correspond to the composition of the market portfolio. Merton (1969, 1971, 1973) shows that long-horizon risk averse investors optimally tilt their portfolios toward those assets whose realized returns are most negatively correlated with unexpected changes in expected returns, because they help hedge their wealth—and consumption—against a deterioration in investment opportunities.

The importance of understanding the optimal value and growth tilts in the portfolios of long-horizon investors is further underscored by the composition of the retail mutual fund universe. According to the CRSP Mutual Fund Database, as of the second quarter of 2005, there were 3797 diversified, domestic equity mutual funds with roughly 2.32 trillion dollars in assets under management.³ Of these funds, 1748 (46%) were classified by CRSP as dedicated growth funds and 1219 (32%) were classified as dedicated value funds, with the remaining 830 (22%) being classified as blend funds. Funds with a dedicated value or growth tilt accounted for 76% of total assets under management (36% growth and 42% value). Thus value and growth tilts are the norm, rather than the exception, in the mutual fund industry that serves the investment needs of most retail investors.

Recent work by Campbell and Vuolteenaho (2004) and others has additionally documented that value and growth stocks differ in their risk characteristics. In particular, the conditional correlation of returns with variables that proxy for time variation in aggregate stock market discount rates is larger for growth stocks than for value stocks, while the conditional correlation of returns with changes in aggregate stock market cash flows is larger for value stocks than for growth stocks. They argue that this should make value stocks riskier than growth stocks from the perspective of a long-horizon risk averse investor, because empirically changes in aggregate stock discount rates are transitory, while changes in aggregate expected cash flows are largely permanent. In fact, they show that an unconditional two-factor model, where one factor captures cash flow risk and the other discount rate risk, can explain the average returns on the Fama and French (1992, 1993, 1996) book-to-market portfolios.

Building on the intuition in Campbell and Vuolteenaho, we compute the optimal portfolio allocation to value and growth of risk-averse investors, and examine how this allocation

³We classify a mutual fund as an *equity* fund if its holdings of cash and common equities account for over 90% of the portfolio value. *Diversified* equity funds exclude sector funds with total net assets under management of 191 billion dollars.

changes across investment horizons. To this end, we model investment opportunities using a vector autoregressive model that includes the returns on growth and value stocks, as well as variables that proxy for expected aggregate stock returns. Additionally, we explore the robustness of these results to the inclusion of other assets, such as T-bills and long-term bonds, in the investment opportunity set while allowing for temporal variation in expected bond excess returns, real interest rates, and inflation.

In related work, Brennan and Xia (2001) and Lynch (2001) also examine optimal dynamic allocations to Fama and French size and book-to-market zero-investment portfolios. However, there are important differences between those papers and our work. Brennan and Xia (2001) ignore time variation in investment opportunities, and focus on the value spread as a “data anomaly” whose existence as a real phenomenon is assessed by the Bayesian investor. Thus their focus is not on long-horizon risk, but on parameter uncertainty and learning.

Our empirical application is closest to Lynch (2001) which explores optimal value and size tilts in the portfolios of long-horizon power utility investors when investment opportunities are time varying. In the paper time variation in investment opportunities is described by the dividend yield on the aggregate stock market and the spread between the long and the short nominal interest rates. The paper uses standard numerical methods to solve the model for a limited set of parameter values and state variables. Our analytical solution allows us to consider a continuous range of parameter values, a richer specification of the state vector, and facilitates the inclusion of additional assets such as long-term bonds in the investment universe. We show that this inclusion results in important qualitative differences in the way long-term investors choose to tilt their equity portfolios.

The organization of the paper is as follows. Section 2 specifies investment opportunities and investor’s preferences, and it states the intertemporal optimization problem. Section 3 solves the model when investors are constrained to follow a constant proportion portfolio strategy. Section 4 solves the problem when investors can dynamically rebalance their portfolios and change portfolio weights in response to changes in investment opportunities. Section 5 applies our method to the empirically relevant problem of constructing an optimal long-term portfolio of value stocks, growth stocks, bonds, and bills given historically estimated return processes. Finally, Section 6 concludes. The Appendix provides a detailed derivation of all the analytical results in the paper.

2 Investment opportunities and investors

We start by outlining our assumptions about the dynamics of the available investment opportunities. We then turn to an analysis of the effect of intertemporal variation in the investment opportunity set on the moments of risky asset returns at long-horizons, and finally, we formalize the investor’s optimization problem.

2.1 Investment opportunities

We consider an economy with multiple assets available for investment, where expected returns and interest rates are time varying. We assume that asset returns and the state variables that characterize time variation in expected returns and interest rates are jointly determined by a first-order linear vector autoregression, or VAR(1):

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

Here \mathbf{z}_{t+1} denotes an $(m \times 1)$ column vector whose elements are the returns on all assets under consideration, and the values of the state variables at time $(t + 1)$. $\mathbf{\Phi}_0$ is a vector of intercepts, and $\mathbf{\Phi}_1$ is a square matrix that stacks together the slope coefficients. Finally, \mathbf{v}_{t+1} is a vector of zero-mean shocks to the realizations of returns and return forecasting variables. We assume these shocks are homoskedastic, and normally distributed:⁴

$$\mathbf{v}_{t+1} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \mathbf{\Sigma}_v). \quad (2)$$

For convenience for our subsequent portfolio analysis, we write the vector \mathbf{z}_{t+1} as

$$\mathbf{z}_{t+1} \equiv \begin{bmatrix} r_{1,t+1} \\ \mathbf{r}_{t+1} - r_{1,t+1} \mathbf{1} \\ \mathbf{s}_{t+1} \end{bmatrix} \equiv \begin{bmatrix} r_{1,t+1} \\ \mathbf{x}_{t+1} \\ \mathbf{s}_{t+1} \end{bmatrix}, \quad (3)$$

where $r_{1,t+1}$ denotes the log real return on the asset that we use as a benchmark in excess return computations, \mathbf{x}_{t+1} is a vector of log excess returns on all other assets 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 that there are $n + 1$ assets, and $m - n - 1$ state variables.

⁴While the simplifying assumption of time invariant risk is perhaps not empirically plausible, it is nonetheless relatively harmless given our focus on long-term portfolio choice decisions. Using a realistically calibrated model of stock return volatility, Chacko and Viceira (2005) argue that the persistence and volatility of risk are not large enough to have a sizable impact on the portfolio decisions of long-term investors, relative to the portfolio decisions of short-horizon investors.

Consistent with our representation of \mathbf{z}_{t+1} in (3), we can write Σ_v as

$$\Sigma_v \equiv \text{Var}_t(\mathbf{v}_{t+1}) = \begin{bmatrix} \sigma_1^2 & \sigma'_{1x} & \sigma'_{1s} \\ \sigma_{1x} & \Sigma_{xx} & \Sigma'_{xs} \\ \sigma_{1s} & \Sigma_{xs} & \Sigma_s \end{bmatrix},$$

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 shocks to 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_{1x} and σ_{1s}), and the covariances of excess returns with shocks to the state variables (Σ_{xs}).

2.2 Long-horizon asset return moments

Despite the seemingly restrictive assumption of homoskedasticity of the VAR shocks, the vector autoregressive specification is able to capture a rich set of dynamics in the moments of long-horizon asset returns. In particular, at horizons exceeding one period asset return predictability generates variation in *per period* risk and expected gross returns (or arithmetic mean returns) across investment horizons, regardless of whether the conditional second moments of the VAR shocks are constant over time or not. We emphasize these implications of asset return predictability because they are useful in understanding horizon effects on portfolio choice.

Consider the conditional variance of K -period log excess returns,

$$\text{Var}_t[\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K}] \equiv \Sigma_{xx}^{(K)},$$

where we have defined $\mathbf{r}_{t \rightarrow t+K} = \sum_{i=1}^K \mathbf{r}_{t+i}$, and $r_{1,t \rightarrow t+K} = \sum_{i=1}^K r_{1,t+i}$. Of course, $\Sigma_{xx}^{(1)}$ is simply the conditional variance of one-period excess returns, Σ_{xx} .

We show in the Appendix that when expected returns are constant—that is, when the slope coefficients in the equations for excess returns in the VAR(1) model are all zero—, $\Sigma_{xx}^{(K)}/K = \Sigma_{xx}$ at all horizons. By contrast, return predictability implies that $\Sigma_{xx}^{(K)}/K$ will generally be different from Σ_{xx} , thus generating a term structure of risk (Campbell and Viceira 2005). Similar considerations apply to the conditional variance of K -period returns on the benchmark asset, which we denote by $(\sigma_1^{(K)})^2$, and the conditional covariance of excess returns with the return on the benchmark asset, which we denote by $\sigma_{1x}^{(K)}$.

Return predictability also generates a term structure of expected gross returns. To see this, note that the log of the unconditional mean gross excess return per period at horizon

K (or the log of the population arithmetic mean return) is related to the unconditional mean log excess return per period at horizon K (the population geometric mean return) as follows:⁵

$$\begin{aligned} \frac{1}{K} \log \mathbb{E} [\exp(\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota)] &= \mathbb{E} [\mathbf{r}_{t+1} - r_{1,t+1} \iota] + \frac{1}{2K} \text{diag} \left(\Sigma_{xx}^{(K)} \right) \\ &+ \frac{1}{2K} \text{Var} [\mathbb{E}_t [\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota]] \end{aligned} \quad (4)$$

Equation (4) implies that the arithmetic mean return is horizon dependent in general, while the geometric mean return is horizon independent. The dependence of the arithmetic mean return on horizon operates through the variance terms, which do not grow linearly with horizon unless returns are not predictable. In the special case of no return predictability, we have that $\Sigma_{xx}^{(K)} = K \Sigma_{xx}$ and $\text{Var} [\mathbb{E}_t [\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota]] = 0$.⁶

Equation (4) gives us strong intuition about the set of investors for whom horizon effects are important. It suggests that investment horizon considerations will be irrelevant to investors who care only about maximizing the geometric mean return on their wealth, while they will be highly relevant to investors for whom the criterion for making portfolio decisions is the arithmetic mean return on their wealth.

Figure 1 gives an empirical illustration of horizon effects on expected returns. This figure plots the annualized geometric mean return (dash-dot line) and annualized arithmetic mean return (solid line) on U.S. stocks and a constant maturity 5-year Treasury bond as a function of investment horizon. The figure considers investment horizons between 1 month and 300 months (or 25 years).⁷ The geometric average return per period of course does not change with the horizon, but the arithmetic mean return per period does change significantly. For U.S. equities, it goes from about 5.3% per year at a 1-quarter horizon to about 4.9% at a 25-year horizon. For U.S. bonds, it decreases from about 1.8% per year to about 1.7% per

⁵This equation follows immediately from applying a standard variance decomposition result:

$$\begin{aligned} \log \mathbb{E} [\exp(\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota)] &= \mathbb{E} [\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota] + \frac{1}{2} \left(\mathbb{E} [\text{Var}_t [\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota]] + \right. \\ &\left. + \text{Var} [\mathbb{E}_t [\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota]] \right) \end{aligned}$$

⁶More generally $\text{Var} [\mathbb{E}_t [\mathbf{r}_{t \rightarrow t+K} - r_{1,t \rightarrow t+K} \iota]]$ equals the elements in the diagonal of

$$\left(\sum_{j=1}^K \Phi_1^j \right)' \text{Var}[\mathbf{z}_t] \left(\sum_{j=1}^K \Phi_1^j \right)$$

corresponding to log excess returns.

⁷This figure is based on a VAR(1) system estimated using postwar monthly data. The VAR includes the same state variables as the VAR we use in our empirical application. See Section 5 for details.

year. The declining average arithmetic return is the direct result of a pattern of decreasing volatility per period of stock and bond returns across investment horizons, which is more pronounced for stocks than for bonds (Campbell and Viceira, 2005).

2.3 Investor's Problem

We consider an investor with initial wealth W_t at time t who chooses a portfolio strategy that maximizes the expected utility of her wealth K periods ahead. At the terminal date, $t + K$, the investor consumes all of the wealth she has accumulated. The investor has isoelastic preferences, with constant coefficient of relative risk aversion γ .

Formally, the investor chooses the sequence of portfolio weights $\boldsymbol{\alpha}_{t+K-\tau}$ between time t and $(t + K - 1)$ such that

$$\left\{ \boldsymbol{\alpha}_{t+K-\tau}^{(\tau)} \right\}_{\tau=K}^{\tau=1} = \arg \max \mathbb{E}_t \left[\frac{1}{1-\gamma} W_{t+K}^{1-\gamma} \right], \quad (5)$$

when $\gamma \neq 1$, and

$$\left\{ \boldsymbol{\alpha}_{t+K-\tau}^{(\tau)} \right\}_{\tau=K}^{\tau=1} = \arg \max \mathbb{E}_t [\log (W_{t+K})].$$

when $\gamma = 1$. Note that we index the sequence of portfolio weight vectors by the time at which they are chosen (subscript) and the time-remaining to the horizon (superscript τ).

Investor's wealth evolves over time as:

$$W_{t+1} = W_t (1 + R_{p,t+1}), \quad (6)$$

where $(1 + R_{p,t+1})$, the gross return on wealth, is given by:

$$\begin{aligned} 1 + R_{p,t+1} &= \sum_{j=1}^n \alpha_{j,t} (R_{j,t+1} - R_{1,t+1}) + (1 + R_{1,t+1}) \\ &= \boldsymbol{\alpha}'_t (\mathbf{R}_{t+1} - R_{1,t+1} \boldsymbol{\iota}) + (1 + R_{1,t+1}), \end{aligned} \quad (7)$$

which is a linear function of the vector of portfolio weights at time t .

Equation (6) implies terminal wealth W_{t+K} is equal to the initial wealth W_t multiplied by the cumulative K -period gross return on wealth, which itself is a function of the sequence of decision variables $\{\boldsymbol{\alpha}_{t+K-\tau}^{(\tau)}\}_{\tau=K}^{\tau=1}$:

$$W_{t+K} = W_t \cdot \prod_{i=1}^K (1 + R_{p,t+i}(\boldsymbol{\alpha}_{t+i-1}^{(K-i+1)})).$$

The preference structure in the model implies that the investor always chooses a portfolio policy such that $(1 + R_{p,t+i}) > 0$.⁸ Thus, along the optimal path,

$$W_{t+K} = W_t \exp \{r_{p,t \rightarrow t+K}\}, \quad (8)$$

where $r_{p,t \rightarrow t+K} = \sum_{i=1}^K r_{p,t+i}$ is the K -period log return on wealth between times t and $t + K$.

Using (8) we can rewrite the objective function (5) as:

$$\left\{ \alpha_{t+K-\tau}^{(\tau)} \right\}_{\tau=K}^{\tau=1} = \arg \max \frac{1}{1-\gamma} \text{E}_t [\exp \{(1-\gamma) r_{p,t \rightarrow t+K}\}], \quad (9)$$

where for simplicity we have dropped the scaling factor $W_t^{1-\gamma}$, which is irrelevant for optimality conditions. Similarly, when $\gamma = 1$, the objective function (9) becomes:

$$\left\{ \alpha_{t+K-\tau}^{(\tau)} \right\}_{\tau=K}^{\tau=1} = \arg \max \text{E}_t [r_{p,t \rightarrow t+K}]. \quad (10)$$

Equation (9) says that a power utility investor with $\gamma \neq 1$ seeks to maximize a power function of the expected long-horizon *gross* return on wealth. By contrast, equation (10) says that a log utility investor seeks to maximize the expected long-horizon *log* return on wealth. Before formally deriving the optimal portfolio policies implied by these two objective functions, we can already say that they will be qualitatively different by simply recalling the properties of long-horizon arithmetic returns and geometric returns shown in Section 2.2.

Finally, following CCV (2003) we approximate the log return on the wealth portfolio (7) as:

$$r_{p,t+1} \approx r_{0,t+1} + \alpha_t' (\mathbf{r}_{t+1} - r_{1,t+1}\mathbf{1}) + \frac{1}{2} \alpha_t' (\sigma_x^2 - \Sigma_{xx} \alpha_t), \quad (11)$$

where $\sigma_x^2 \equiv \text{diag}(\Sigma_{xx})$ is the vector consisting of the diagonal elements of Σ_{xx} , the variances of log excess returns. Equation (11) is an approximation which becomes increasingly accurate as the frequency of portfolio rebalancing increases, and it is exact in the continuous time limit.

⁸To see this, note that a zero one-period gross return on wealth at any date implies zero wealth and consumption at time $t + K$, which in turn implies that marginal utility of consumption approaches infinity. This is a state of the world the investor will surely avoid.

3 Optimal constant proportion portfolio strategies

Before we solve the optimal dynamic portfolio choice problem described in Section 2, it is useful to consider a simpler case that constrains a K -period investor to choose at the beginning of her investment horizon a portfolio strategy with constant portfolio weights:

$$\left\{ \boldsymbol{\alpha}_{t+K-\tau}^{(\tau)} \right\}_{\tau=K}^{\tau=1} = \bar{\boldsymbol{\alpha}}_t^{(K)}, \quad (12)$$

where $\bar{\boldsymbol{\alpha}}_t^{(K)}$ denotes the vector of constrained constant portfolio weights optimally chosen at time t by an investor with a K -period investment horizon.

Solving for the optimal constrained policy (12) and studying its properties is useful for several reasons. First, many institutional and individual investors frequently adopt long-horizon asset allocation policies that call for rebalancing the portfolio to a fixed vector of portfolio weights. Thus the constrained constant proportion strategy provides a benchmark of practical relevance for the unconstrained dynamic rebalancing strategy resulting from (9). Second, the constrained solution is useful in building intuition about the unconstrained solution, which we develop in Section 4.⁹ Third, the solution to the constrained problem gives interesting insights on the connection between horizon effects on portfolio choice and the different behavior of geometric mean returns and arithmetic mean returns across investment horizons which we have pointed out in Section 2.

Solving for the optimal constant proportion strategy requires computing the expectation in equation (9). Given our assumptions about the conditional distribution of returns, the continuous-time approximation to the log return on wealth (11) implies that one-period log portfolio returns are conditionally normal. This in turn implies that the cumulative long-horizon log return on wealth $r_{p,t \rightarrow t+K}$ is also conditionally normal *when portfolio weights are constant over time*. Thus the long-horizon gross return on wealth $\exp\{r_{p,t \rightarrow t+K}\}$ is conditionally lognormal and we can compute the expectation in (9) as:

$$\mathbb{E}_t [\exp \{(1 - \gamma) r_{p,t \rightarrow t+K}\}] = \exp \left\{ (1 - \gamma) \mathbb{E}_t [r_{p,t \rightarrow t+K}] + \frac{1}{2} (1 - \gamma)^2 \text{Var}_t [r_{p,t \rightarrow t+K}] \right\}, \quad (13)$$

⁹It is important to note that the K -horizon constant proportion strategy is not the limiting case of the K -horizon dynamically rebalanced strategy as the rebalancing frequency approaches 0 (i.e., no rebalancing). This limit is instead a buy-and-hold strategy where investors do not rebalance their portfolios at all over their investment horizon. Campbell and Viceira (2005) explore this strategy in a mean-variance framework. By contrast, in the K -horizon constant proportion strategy investors rebalance their portfolios every period (i.e. at frequency $1/K$).

where $E_t[r_{p,t \rightarrow t+K}]$ and $\text{Var}_t[r_{p,t \rightarrow t+K}]$ are the conditional moments of the K -period log return on wealth evaluated under (12). Closed-form expressions for these moments are provided in the Appendix.

Substitution of equation (13) into (9) gives the objective function for a long-horizon investor with a coefficient of relative risk aversion different from unity. This objective function is equivalent to

$$\bar{\alpha}_t^{(K)} = \arg \max \left\{ \left(E_t[r_{p,t \rightarrow t+K}] + \frac{1}{2} \text{Var}_t[r_{p,t \rightarrow t+K}] \right) - \frac{\gamma}{2} \text{Var}_t[r_{p,t \rightarrow t+K}] \right\}. \quad (14)$$

Thus a constant-proportion power utility investor with $\gamma \neq 1$ seeks to maximize the expected long-horizon gross return on wealth subject to a constraint on long-horizon wealth return volatility. By contrast, the objective function (10) for a log utility investor shows that this investor seeks to maximize the expected long-horizon *log* return on wealth, regardless of volatility.

Solving for $\bar{\alpha}_t^{(K)}$ in (14) leads to:

$$\bar{\alpha}_t^{(K)} = \left(\Sigma_{xx} - (1 - \gamma) \frac{1}{K} \Sigma_{xx}^{(K)} \right)^{-1} \left(\frac{1}{K} E_t[\mathbf{r}_{t \rightarrow t+K} - r_{0,t \rightarrow t+K} \mathbf{1}] + \frac{1}{2} \boldsymbol{\sigma}_x^2 + (1 - \gamma) \frac{1}{K} \boldsymbol{\sigma}_{1x}^{(K)} \right). \quad (15)$$

For log utility investors ($\gamma = 1$), equation (15) reduces to:

$$\bar{\alpha}_t^{(K)} = \Sigma_{xx}^{-1} \left(\frac{1}{K} E_t[\mathbf{r}_{t \rightarrow t+K} - r_{0,t \rightarrow t+K} \mathbf{1}] + \frac{1}{2} \boldsymbol{\sigma}_x^2 \right), \quad (16)$$

which is similar to the well-known mean-variance tangency portfolio, except that it is based on the conditional expectation of the K -period excess returns, instead of the one-period expected excess return.

Although the optimal constant proportion portfolio strategy (15) is time-invariant by construction, it obviously depends on the value of state vector \mathbf{z}_t at time t through the conditional expectation of the K -period log portfolio return. Thus to focus on horizon effects it is convenient to analyze the portfolio allocations at their unconditional means, $E[\boldsymbol{\alpha}]$. Taking unconditional expectations on both sides of (15) and (16) it is easy to see that the mean allocation of the log utility investor is independent of horizon, while the mean allocation for all other power utility investors does depend on investment horizon, through the terms $\Sigma_{xx}^{(K)}/K$ and $\boldsymbol{\sigma}_{1x}^{(K)}/K$.¹⁰

¹⁰Of course, when expected returns are constant we have that $\Sigma_{xx}^{(T)}/T = \Sigma_{xx}$ and $\boldsymbol{\sigma}_{0x}^{(T)}/T = \boldsymbol{\sigma}_{0x}$, and $E[\boldsymbol{\alpha}^{(T)}] = E[\boldsymbol{\alpha}^{(1)}]$ for all T .

Equation (15) shows that long-horizon investors with $\gamma \neq 1$ choose portfolio weights which vary inversely with a weighted average of the variance-covariance matrix of excess returns at short- and long-horizons.¹¹ The weight on long-horizon return variance is increasing in the distance between one and the value of the coefficient of relative risk aversion. Thus log utility investors put no weight on long-horizon return variance, as (16) shows, while increasingly risk averse investors put more weight on this variance. In the limit, as $\gamma \rightarrow \infty$, the optimal mean portfolio allocation (15) approaches:

$$\bar{\alpha}_t^{(K)} \rightarrow - \left(\Sigma_{xx}^{(K)} \right)^{-1} \sigma_{1x}^{(K)}$$

which is the global minimum variance (GMV) portfolio at horizon K .

4 A general recursive solution

We now solve for the unconstrained, dynamic rebalancing strategy proceeding by standard backwards recursion. We first derive the portfolio rule in the last period (the base case for the policy function recursion) and the associated value function (the base case for the value function recursion). We then solve the problem for the period preceding the last portfolio choice date as a function of the value and policy function coefficients from the terminal period. This enables us to isolate the recursive relationship linking the policy function and value function recursions for two adjacent periods. By iterating this relationship we arrive at the solution to the general multi-period portfolio choice problem with dynamic rebalancing.

Our solution possesses a variety of attractive features. First, it flexibly accommodates any number of risky assets and state variables. Second, it is exact in the limit when the investor can rebalance her portfolio continuously, since the loglinear approximation (11) is exact in continuous time. And lastly, in the special cases when there is only one period remaining or returns are not predictable, our solution simplifies to the well known myopic portfolio choice rule.

¹¹Campbell and Viceira (2005) show that the optimal portfolio of a long-horizon, buy-and-hold investor puts full weight on $\Sigma_{xx}^{(T)}/T$ and no weight on Σ_{xx} , regardless of risk aversion.

4.1 Optimal portfolio policy when remaining horizon is one period

Equation (9) implies that the objective for an investor with a remaining horizon of one period is to choose $\boldsymbol{\alpha}_{t+K-1}^{(1)}$ so that:

$$\boldsymbol{\alpha}_{t+K-1}^{(1)} = \arg \max \frac{1}{1-\gamma} \mathbf{E}_{t+K-1} [\exp \{(1-\gamma) r_{p,t+K}\}]. \quad (17)$$

The one-period problem is formally analogous to the constant proportion strategy problem of Section 3, since under the approximation (11) and the distribution assumption (2), $r_{p,t+K}$ is conditionally lognormal. Thus specializing $K = 1$ in (15) we obtain the solution to the optimization problem (17):

$$\boldsymbol{\alpha}_{t+K-1}^{(1)} = \frac{1}{\gamma} \boldsymbol{\Sigma}_{xx}^{-1} \left(\mathbf{E}_{t+K-1} [\mathbf{r}_{t+K} - r_{1,t+K} \boldsymbol{\iota}] + \frac{1}{2} \boldsymbol{\sigma}_x^2 + (1-\gamma) \boldsymbol{\sigma}_{1x} \right), \quad (18)$$

where

$$\mathbf{E}_{t+K-1} [\mathbf{r}_{t+K} - r_{0,t+K} \boldsymbol{\iota}] = \mathbf{H}_x (\boldsymbol{\Phi}_0 + \boldsymbol{\Phi}_1 \mathbf{z}_{t+K-1}), \quad (19)$$

and \mathbf{H}_x is a matrix operator that selects the rows corresponding to the vector of excess returns \mathbf{x} from the target matrix. Thus the solution (18) implicitly defines an affine function of the state vector \mathbf{z}_{t+K-1} .

Equation (18) is the well-known ‘‘myopic’’ or one-period mean-variance efficient portfolio rule. The optimal myopic portfolio (18) combines the tangency portfolio and the global minimum variance portfolio of the mean-variance efficient frontier generated by one-period expected returns and the conditional variance-covariance matrix of one-period returns. The tangency portfolio is:

$$\boldsymbol{\Sigma}_{xx}^{-1} \mathbf{E}_{t+K-1} [\mathbf{r}_{t+K} - r_{1,t+K} \boldsymbol{\iota}] + \frac{1}{2} \boldsymbol{\sigma}_x^2. \quad (20)$$

This portfolio depends on expected returns and the variance-covariance matrix of returns. In our model, expected returns are time-varying, causing this portfolio to change with the investment opportunities. The global minimum variance (GMV) portfolio is

$$-\boldsymbol{\Sigma}_{xx}^{-1} \boldsymbol{\sigma}_{1x}, \quad (21)$$

and depends only on the variance-covariance structure of returns. Our assumption of constant variances and covariances implies that the single-period GMV portfolio does not change with investment opportunities.

Investors combine these two portfolios using weights $1/\gamma$ and $(1 - 1/\gamma)$, respectively. Log utility investors (investors with unit coefficient of relative risk aversion γ) hold only the tangency portfolio, while highly risk averse investors (investors for whom $\gamma \rightarrow \infty$) hold only the GMV portfolio. Other investors hold a mixture of both.

4.2 Value function when remaining horizon is one period

Since $r_{p,t+K}$ is conditionally lognormal, we can write the objective function (17) as

$$\mathbb{E}_{t+K-1} [\exp \{(1-\gamma)r_{p,t+K}\}] \propto \exp \left\{ (1-\gamma) \mathbb{E}_{t+K-1} [r_{p,t+K}] + \frac{1}{2} (1-\gamma) \text{Var}_{t+K-1} [r_{p,t+K}] \right\}$$

times the scaling factor $1/(1-\gamma)$. Thus the value function at time $(t+K-1)$ depends on the expected log return on wealth and its variance. Substitution of the optimal portfolio rule (18) into the equation for the log return on wealth (11) leads to expressions for the expected log return on wealth and its variance which are both quadratic functions of the \mathbf{z}_{t+K-1} vector. This is intuitive, since the expected log return on wealth depends on the product of $\boldsymbol{\alpha}_{t+K-1}^{(1)}$ and the expected return on wealth, both of which are linear in \mathbf{z}_{t+K-1} ; similarly, the conditional variance of the log return on wealth depends quadratically on $\boldsymbol{\alpha}_{t+K-1}^{(1)}$, which is itself a linear function of \mathbf{z}_{t+K-1} . Therefore, the expectation in the value function at time $(t+K-1)$ is itself an exponential quadratic polynomial of \mathbf{z}_{t+K-1} :

$$\mathbb{E}_{t+K-1} [\exp \{(1-\gamma)r_{p,t+K}\}] = \exp \left\{ (1-\gamma) \left(B_0^{(1)} + B_1^{(1)} \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} B_2^{(1)} \mathbf{z}_{t+K-1} \right) \right\} \quad (22)$$

where $B_0^{(1)}$, $B_1^{(1)}$, and $B_2^{(1)}$ are given in the Appendix.

4.3 Optimal portfolio policy and value function when remaining horizon is two periods

We now proceed to compute the optimal portfolio policy and the value function at time $(t+K-2)$. When the remaining horizon is two periods, the investor's objective function is

$$\max_{\boldsymbol{\alpha}_{t+K-2}^{(2)}, \boldsymbol{\alpha}_{t+K-1}^{(1)}} \frac{1}{1-\gamma} \mathbb{E}_{t+K-2} [\exp \{(1-\gamma)(r_{p,t+K-1} + r_{p,t+K})\}] \quad (23)$$

which, using the law of iterated expectations and equation (22), we can further rewrite as

$$\max_{\boldsymbol{\alpha}_{t+K-2}^{(2)}} \frac{1}{1-\gamma} \mathbb{E}_{t+K-2} \left[\exp \left\{ (1-\gamma) \left(r_{p,t+K-1} + B_1^{(1)} \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} B_2^{(1)} \mathbf{z}_{t+K-1} \right) \right\} \right]. \quad (24)$$

In order to compute the optimal portfolio policy and the value function at $(t+K-2)$, we need to evaluate the expectation (24). Note that the last two terms inside the expectation define an affine-quadratic function of \mathbf{z}_{t+K-1} , and that equation (11) implies that $r_{p,t+K-1}$

is an affine function of \mathbf{z}_{t+K-1} . Thus the term inside the expectation is an exponential quadratic polynomial function of the vector of state variables \mathbf{z}_{t+K-1} . We can evaluate this expectation in closed form using standard results on the expectation of an exponential quadratic polynomial of normal variables. The Appendix provides an analytical expression for (24), and some additional simplifications applicable in the continuous-time limit.

The analytical evaluation of the expectation (24) results in an objective function¹² whose first order condition implies an optimal portfolio policy which, similar to the optimal one-period portfolio policy, is also an affine function of the state vector \mathbf{z}_{t+K-2} . It is important however to note that the coefficients of this function will, in general, be different from the coefficients of the state vector \mathbf{z}_{t+K-1} in the one-period solution. They differ in qualitatively important ways that capture the fact that the optimal portfolio rule is not necessarily myopic when the remaining investment horizon is longer than one period—and the agent anticipates further opportunities for portfolio rebalancing in the face of changing investment opportunities. We defer the discussion of these differences until we present the general solution at any remaining horizon τ in the next section.

Similarly, substitution of the optimal portfolio policy $\boldsymbol{\alpha}_{t+K-2}^{(2)}$ back into the objective function leads to a value function at $(t + K - 2)$ which has the same functional form as the value function (22) at $(t + K - 1)$, but with coefficients $B_0^{(2)}$, $B_1^{(2)}$, and $B_2^{(2)}$ which are not necessarily equal to the coefficients of the one-period value function. The Appendix provides expressions for these coefficients.

4.4 General recursive solution and its properties

The results for the one-period horizon case and the two-period horizon case implicitly define the recursion generating the portfolio rule at any horizon. The one-period horizon solution represents the base case for the recursive solution, while the two-period horizon solution effectively relates the portfolio allocation in two adjacent time periods.

With τ time periods remaining the optimal portfolio rule is given by

$$\begin{aligned} \boldsymbol{\alpha}_{t+K-\tau}^{(\tau)} &= \frac{1}{\gamma} \boldsymbol{\Sigma}_{xx}^{-1} \left(\mathbb{E}_{t+K-\tau} [\mathbf{r}_{t+K-\tau+1} - r_{1,t+K-\tau+1} \mathbf{1}] + \frac{1}{2} \boldsymbol{\sigma}_x^2 + (1 - \gamma) \boldsymbol{\sigma}_{1x} \right) \\ &\quad - \left(1 - \frac{1}{\gamma} \right) \boldsymbol{\Sigma}_{xx}^{-1} \boldsymbol{\Sigma}_x \left(B_1^{(\tau-1)'} + \left(B_2^{(\tau-1)} + B_2^{(\tau-1)'} \right) \mathbb{E}_{t+K-\tau} [\mathbf{z}_{t+K-\tau+1}] \right), \end{aligned} \quad (25)$$

¹²The objective function is itself an exponential quadratic polynomial function of \mathbf{z}_{t+K-2} whose coefficients depend on $\boldsymbol{\alpha}_{t+K-2}^{(2)}$, the decision variable. Viewed as a function of $\boldsymbol{\alpha}_{t+K-2}^{(2)}$, the objective function is also an exponential quadratic polynomial function of $\boldsymbol{\alpha}_{t+K-2}^{(2)}$.

where $\Sigma_x = [\sigma_{1x} \quad \Sigma_{xx} \quad \Sigma'_{xs}]$.

Equation (25) shows that the optimal portfolio demand at horizon $\tau > 1$ has two components. The first component is identical to the one-period myopic portfolio demand (18), and it is independent of investment horizon. The second component reflects an additional intertemporal hedging portfolio demand for risky assets which is absent at $\tau = 1$, when investors have only one period to go before liquidating their assets and consuming their wealth. When $\tau > 1$, investors have more time and more opportunities for rebalancing before liquidating their assets than when $\tau = 1$, which leads them to care about future changes in investment opportunities. Risk averse investors might want to tilt their portfolios toward assets that protect their wealth from adverse changes in investment opportunities (Merton 1969, 1971, 1973).

The intertemporal hedging component of total portfolio demand takes the form of an affine function of the vector of state variables:

$$\begin{aligned} & - \left(1 - \frac{1}{\gamma} \right) \Sigma_{xx}^{-1} \Sigma_x \left(B_1^{(\tau-1)'} + \left(B_2^{(\tau-1)} + B_2^{(\tau-1)'} \right) \mathbf{E}_{t+K-\tau} [\mathbf{z}_{t+K-\tau+1}] \right) \\ = & - \left(1 - \frac{1}{\gamma} \right) \Sigma_{xx}^{-1} \Sigma_x \left(B_1^{(\tau-1)'} + \left(B_2^{(\tau-1)} + B_2^{(\tau-1)'} \right) (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-\tau}) \right). \end{aligned}$$

The coefficients of this function depend on the coefficient of relative risk aversion, the coefficients of the VAR system, and the coefficients of the value function, which are themselves functions of the coefficients of the VAR system and the remaining investment horizon τ . Thus it is through the intertemporal hedging component that the investment horizon τ affects the portfolio demand of the dynamically rebalancing investor.

Since the myopic component of total portfolio demand is also an affine function of the vector of state variables $\mathbf{z}_{t+K-\tau}$, we can rewrite total portfolio demand $\alpha_{t+K-\tau}^{(\tau)}$ as an affine function of the vector of state variables:

$$\alpha_{t+K-\tau}^{(\tau)} = A_0^{(\tau)} + A_1^{(\tau)} \mathbf{z}_{t+K-\tau}, \quad (26)$$

where the expressions for $A_0^{(\tau)}$ and $A_1^{(\tau)}$ are obvious from (25). The dynamic consistency of the policy function ensures that the coefficient matrices, $A_0^{(\tau)}$ and $A_1^{(\tau)}$, depend *only* on the time remaining to the terminal horizon date, $t + K$, but are independent of time itself. Consequently, we index these coefficients by the time remaining to the consumption date.

It is also important to note that $A_0^{(\tau)}$ and $A_1^{(\tau)}$ only depend on the remaining investment horizon τ via the intertemporal hedging component of total demand. Thus while both the myopic component and the intertemporal hedging component of total portfolio demand

make this demand state dependent, the dependence of total portfolio demand on horizon is exclusively driven by intertemporal hedging considerations.

It is useful to compare the optimal portfolio rule with dynamic rebalancing with the optimal constant proportion strategy. Equation (25) and equation (15) show that both rules have similar functional form, but there are some important differences that reflect the fact that the rebalancing investor can dynamically change her portfolio in response to changes in investment opportunities, which in turn leads to an effective shortening of the investment horizon. Thus while the constant proportion strategy depends on the long-horizon first and second moments of returns, the dynamically rebalanced strategy depends only on single-period first and second moments of returns. Horizon considerations enter the dynamic rebalancing strategy only through the intertemporal hedging component, which is absent in the constant proportion strategy.

It is important to check that the optimal dynamically rebalanced portfolio policy converges to well-known solutions in certain limiting cases. In the Appendix we show that the optimal dynamic policy (25) reduces to the myopic solution at all horizons only when investors have log utility ($\gamma \rightarrow 1$), and when investment opportunities are constant ($\mathbf{H}_x \Phi_1 = 0$). In those cases we have that $A_0^{(\tau)} = A_0^{(1)}$ and $A_1^{(\tau)} = A_1^{(1)}$ for all τ .

We also show in the Appendix that as we consider increasingly risk averse investors (i.e., as $\gamma \rightarrow \infty$), the optimal portfolio policy becomes decoupled from the intercept vector, Φ_0 , of the VAR(1). That is, the least-risky portfolio from the perspective of a long-horizon dynamic rebalancing investor is independent of the vector of unconditional mean returns. This portfolio converges to the one-period GMV when investment opportunities are constant.

Finally, the value function with τ periods remaining is still an exponential quadratic function of the state vector:

$$\begin{aligned} & \max_{\alpha} \frac{1}{1-\gamma} \mathbb{E}_{t+K-\tau} \left[\exp \left\{ (1-\gamma) \sum_{i=1}^{\tau} r_{p,t+K-\tau+i} \left(\boldsymbol{\alpha}_{t+K+i-(\tau+1)}^{(\tau+1-i)} \right) \right\} \right] \\ &= \frac{1}{1-\gamma} \exp \left\{ (1-\gamma) \left(B_0^{(\tau)} + B_1^{(\tau)} \mathbf{z}_{t+K-\tau} + \mathbf{z}'_{t+K-\tau} B_2^{(\tau)} \mathbf{z}_{t+K-\tau} \right) \right\}. \end{aligned} \quad (27)$$

The Appendix provides a detailed derivation of all these results, as well as expressions for the coefficients of the optimal portfolio policy and the value function.¹³

¹³In brief, we show in the Appendix that the policy and value function coefficients satisfy a linear recursive relation where the $A_i^{(\tau)}$ coefficients depend linearly on the $B_i^{(\tau-1)}$ value function coefficients, and the $B_i^{(\tau)}$ coefficients depend linearly on both the $B_i^{(\tau-1)}$ and the $A_i^{(\tau-1)}$ coefficients. The parameters of this linear recursion are nonlinear functions of the parameters of the VAR(1) system, and the coefficient of relative risk aversion. These expressions are algebraically involved but trivial to program, allowing for the examination

5 Optimal growth and value investing

The empirical analysis of optimal dynamic portfolio choice with time-varying investment opportunities has focused primarily on the choice between a well-diversified portfolio of equities representing the market, cash, and - in some cases - long-term bonds. Although this setup allows for the analysis of horizon effects in the allocation to equities relative to cash or bonds, it is not designed to yield insights into horizon effects in the composition of the optimal equity portfolio. Investors might also want to optimally change the composition of their equity portfolio in response to changes in investment opportunities, for example, if the covariation of equity returns with state variables is not homogeneous across all types of equities.

A number of papers in empirical finance have documented the existence of predictability (or “mean reversion”) in aggregate stock market returns. This finding suggests that either aggregate stock cash flows or aggregate stock market discount rates, or a combination of both, must vary in a predictable manner. Indeed, empirical analysis indicates that time variation in expected aggregate stock returns appears to be driven primarily by transitory predictable movements in aggregate stock discount rates; changes in aggregate stock cash flows are largely unpredictable.

In a recent paper, Campbell and Vuolteenaho (2004) note that one can decompose the covariance of the unexpected return on any stock with the unexpected return on the stock market into the covariance of the return with shocks to aggregate stock cash flows (“stock market cash flow news”) and the covariance of the return with shocks to aggregate stock discount rates (“stock market discount rate news”). They show that empirically unexpected returns on value stocks covary more closely with aggregate cash flow news than with discount rate news, and that this pattern is reversed for the unexpected return on growth stocks. They argue that their empirical finding implies that value stocks are riskier than growth stocks from the perspective of a risk-averse, long-horizon investor who holds the market portfolio, because aggregate cash flow shocks appear to be permanent, while aggregate discount rate shocks appear to be transitory. Using the first order optimality conditions of an infinitely lived investor who holds the market portfolio, they estimate the coefficient of relative risk aversion which would suffice to deter this investor from tilting her portfolio toward value stocks, and find that this coefficient is large for the period 1963 through 2001.

The flexible framework for the analysis of dynamic portfolio choice problems we have developed in Sections 2 through 4 is ideal to investigate value and growth tilts in equity

of portfolio choice problems involving an arbitrary number of assets and state variables. MATLAB routines which execute the policy and value function recursions are available on the authors’ websites.

portfolios systematically. To this end, we consider an empirical specification of our dynamic portfolio choice problem in which investors can invest in two equity portfolios, a portfolio of value stocks and a portfolio of growth stocks. We consider two types of investors. The first type of investor can only invest in these two equity portfolios. This is the type of investor implicit in most representative investor models of equilibrium asset prices, which assume that bonds are in zero net supply. The second type of investor is an investor who can also invest in cash and bonds in addition to value and growth stocks. In both cases we explore optimal value and growth tilts across investment horizons and across varying levels of risk aversion.

5.1 Investment opportunities

5.1.1 Assets, state variables, and data

Following our theoretical framework, we model the dynamics of investment opportunities as a first-order VAR system. As we have already noted, we consider two sets of investable assets and estimate a companion VAR system for each of these investment sets. The first set is comprised only of equities, and consequently, we refer to this set as the “equity-only case.” In this scenario the investor chooses between a value-weighted portfolio of growth stocks and a complementary value-weighted portfolio of value stocks. The value of the two portfolios adds up to the aggregate stock market portfolio. The companion VAR system includes the log real return on the growth stock portfolio (labelled G in tables), the log return on the value stock portfolio (V) in excess of the log return on the growth portfolio ($V-G$), and a set of common state variables which we describe below.¹⁴

We construct the value and growth portfolios using data on six stock portfolios sorted by the ratio of book value of equity to market value of equity (BM) and market capitalization, available from Professor Ken French’s website and based on raw data from CRSP and COMPUSTAT. We combine the BM and size sorted portfolios into three BM portfolios. We build then V as a value-weighted portfolio that includes the portfolio of stocks with the lowest BM ratios and half the portfolio of stocks with medium BM ratios. G has the complementary composition. Figure 2 plots the share of total stock market value of these portfolios over time. On average the value portfolio represents 30% of total market capitalization, and the growth portfolio represents the remaining 70%. Growth represents

¹⁴We consider $V-G$ instead of V and G separately for consistency with the VAR formulation in our portfolio choice model, which assumes that one of the assets in the investment opportunity set acts as a benchmark asset over which we measure excess returns on all other assets. Since this VAR includes only equity portfolios, the benchmark asset must be one of them. We have chosen the growth portfolio as the benchmark asset, but of course this choice is inconsequential to the portfolio choice results.

more than 80% of total market capitalization in three episodes, the early 1930's, the mid-1970's and the end of the 1990's. By contrast, the largest market share of the value portfolio occurs in the late 1940's, late 1960's and in the mid 1980's.

The second investment set adds cash and long-term Treasury bonds to the two equity portfolios, leading us to refer to it as the "equities-and-bonds case." The companion VAR system includes the log excess return on the value portfolio, the log excess return on the growth portfolio, the log excess return on a constant maturity 5-year Treasury bond (B5), the ex-post real rate of return on a 30-day Treasury bill, and the same set of state variables as in the first system. Excess returns are computed using the 30-day Treasury bill as the benchmark asset.

The common set of state variables includes variables known to forecast aggregate stock excess returns, bond excess returns, interest rates, and inflation. The first of these variables is the price-earnings ratio (PE) on the S&P 500, which forecasts aggregate stock returns negatively at long horizons (Campbell and Shiller 1988, 1998, 2005).¹⁵ The rest of the state variables are related to the term structure of interest rates and inflation. We include the short-term nominal interest rate (t30_YIELD), which forecasts aggregate stock returns negatively (Fama and Schwert 1977, Campbell 1987, Glosten et al. 1993); the yield spread (YSPR), which forecasts bond excess returns positively (Fama and Bliss 1987, Fama and French 1989, Campbell and Shiller 1991, Campbell, Chan and Viceira 2003, Campbell and Viceira 2005); and the ex-post real rate of return on a 30-day Treasury bill (t30_REALRET). Note that the ex-post real rate plays a dual role as the real return on an investable asset (Treasury Bills) and as an additional state variable which, together with the nominal short-term interest rate and the yield spread, allow the VAR system to capture the dynamics of inflation and real interest rates.¹⁶

Our empirical measure of PE is the value of the S&P 500 portfolio divided by the ten-year trailing moving average of aggregate earnings on the S&P 500 companies, which we obtain from CRSP and Campbell and Vuolteenaho (2004). The data source for bond returns, interest rates and inflation is CRSP. The return on bonds is the log return on a constant

¹⁵An alternative variable that captures similar information in expected aggregate stock returns is the dividend-price ratio. This ratio forecasts future stock returns positively (Campbell and Shiller 1988, Fama and French 1989, Hodrick 1992, Goetzmann and Jorion 1993). Brandt (1999), Campbell and Viceira (1999, 2005), Campbell, Chan, and Viceira (2003) and others use this variable in empirically calibrated models of portfolio choice with time-varying expected stock returns.

¹⁶In their study of the cross-sectional pricing of value and growth stocks, Campbell and Vuolteenaho (2004) consider an additional stock market variable. This variable is the small-stock value spread (VS), which is known to forecast aggregate stock returns negatively (Eleswarapu and Reinganum 2004, Brennan, Wang and Xia 2004, Campbell and Vuolteenaho 2004). The inclusion of this variable does not make any difference to our results, so we have excluded it from our analysis in the interest of parsimony.

maturity 5-year Treasury bond from the CRSP US Treasury and Inflation database. The nominal short rate is the log yield on a 30-day Treasury bill from CRSP. The yield spread is the difference between the log yield on a five-year discount bond from the CRSP Fama-Bliss files, and the log yield on the 30-day Treasury bill. Finally, we use the CPI inflation series in the CRSP US Treasury and Inflation database to construct the ex-post real short-term interest rate and the real return on the growth stock portfolio. We provide full details of the variable definitions and construction in Table 1.

Because we do not observe the relations between state variables and asset returns, we estimate both VAR systems using monthly data for the period December 1952 through December 2003. We restrict our sample to the post-1952 period because the Federal Reserve kept short-term nominal rates essentially fixed before the Treasury Accord of 1952, making it impossible to capture interest rate dynamics using the pre-1953 data series. In our subsequent portfolio choice exercise, we also assume that investors take the VAR parameter estimates at face value, ignoring estimation uncertainty.

5.1.2 VAR estimates

Table 2 presents descriptive statistics of the variables included in the VAR system. This table provides a clear illustration of the empirical regularity known as “value premium.” While the value stock portfolio, the growth stock portfolio and the aggregate stock market portfolio have almost identical short-term return volatility, the average return on the value stock portfolio is significantly higher than the average return on the growth stock portfolio and the aggregate stock portfolio. The average spread between the returns on value stocks and the returns on growth stocks is about 2.45% per year. This spread however exhibits non-trivial variation over time and has an annualized standard deviation of nearly 7%. Overall, the Sharpe ratio on the value portfolio is 0.61, which is about 47% larger than the Sharpe ratio on the growth portfolio. Thus from a purely myopic perspective, the ex-post performance of the value portfolio suggests that it represents a more attractive investment opportunity than the growth portfolio.

Table 3 presents estimates for the equity-only VAR system. The table has two panels. Panel A reports the slope coefficient estimates with heteroskedasticity and autocorrelation consistent t-statistics below in parenthesis, and bootstrapped 95% confidence intervals in brackets. The bootstrap estimates are generated from sample paths simulated under the null hypothesis that the estimated VAR model represents the true data generating process. The rightmost column in the panel reports the R^2 for each equation, and the p-value of the F-statistic of the overall significance of the slopes in each equation. Panel B reports the percentage standard deviation of the innovations to each equation (on the main diagonal)

and the cross-correlations of the innovations (off the main diagonal).

Panel A in Table 3 shows that own lagged returns forecast returns positively, though only the coefficient on the lagged return on V-G is statistically significant. PE, the nominal short rate, and the ex-post real short rate are all highly significant in the forecasting equation for the real return on the growth stock portfolio. Both PE and the nominal short interest rate forecast the real return on growth stocks negatively, and the ex-post real rate forecasts this return positively. By contrast, none of the state variables is significant in the forecasting equation for the return on V-G. This implies that these variables forecast the return on the value stock portfolio with coefficients which are not statistically different from the coefficients in the equation for the return on the growth stock portfolio. Table 4, which considers the return on the value stock portfolio separately from the return on the growth stock portfolio, confirms this result. In the case of PE, even the point estimate of its coefficient in the V-G equation is essentially zero.

These results show that variables known to forecast aggregate stock market returns also forecast both the returns of growth stocks and the returns on value stocks separately.¹⁷ Of course, that this holds for the growth portfolio is not surprising, since this portfolio represents about 70% of total market capitalization. It is perhaps less obvious that these variables would also forecast the return on the relatively small value stock portfolio.

Panel A in Table 3 also indicates that the state variables are all well described as persistent autoregressive processes. Except for the ex-post real rate, the autoregressive coefficients for all the state variables are above .90 and, in some cases, above .99. This raises the question of whether the processes for the state variables are unit-root processes rather than highly persistent, but ultimately stationary, processes. Recent research (Campbell and Yogo, 2005) supports the second conclusion, which we adopt in our portfolio choice exercise. Finally, it is interesting to note that there exist some cross-forecasting effects in the equations for the state variables. The ex-post real short rate and the yield spread forecast PE positively, while the nominal short rate forecasts the ex-post real short rate positively, and PE negatively.

As the results in Sections 2-4 show, the optimal portfolio allocations depend not only on the slope coefficients in the VAR, but also on the variance-covariance structure of the innovations to the VAR variables. Panel B in Table 3 presents this matrix for the equity-only VAR system. Two facts stand out from this panel.

¹⁷We have estimated a VAR system for the real return on the market portfolio with the same state variables. Our estimate confirms that these three state variables forecast aggregate stock market returns in the same direction and with similar statistical significance in our sample period.

First, the correlation between the unexpected return on growth stocks and the unexpected return on V-G is about -40% . The negative of this correlation determines the sign of the one-period GMV portfolio allocation to value stocks in the equity-only case—see equation (21). Its negative sign implies that this portfolio loads positively on value stocks.

Second, shocks to PE are highly positively correlated with unexpected returns on the growth stock portfolio, and negatively correlated with unexpected returns in V-G. This implies that innovations to PE are more correlated with unexpected returns on the growth portfolio than with unexpected returns on the value portfolio. Since innovations to PE proxy for innovations to expected returns on the aggregate stock market portfolio, our results suggest—consistent with the evidence in Campbell and Vuolteenaho (2004)—that returns on growth stocks covary more with stock market discount rate news than the returns on value stocks. We validate this conclusion by bootstrapping the distribution of the covariance estimator under the null hypothesis of an equal covariance between innovations to value and growth and PE. Our simulation reveals that the historical estimate of the covariance between shocks to the value-growth spread and PE is sufficient to reject the null hypothesis at the 0.01%-level (the maximal theoretical level possible under a simulation with 10,000 paths).

Table 4 presents estimates for the equities-and-bonds VAR system. We present constrained estimates in which only the lagged values of the state variables forecast future returns. We have chosen to constrain this system because lagged returns, particularly lagged bond returns, forecast next period’s returns at monthly frequencies positively. This short-lived “momentum” effect induces short-horizon effects in portfolio allocations unrelated to the effects of persistent changes in investment opportunities, which are well captured by our state variables. Because our primary interest is in long-horizon effects in portfolio allocations, we choose to present constrained estimates for the equities-and-bonds VAR system, and portfolio allocations based on these constrained estimates.¹⁸

The estimation results for the equities-and-bonds VAR system are consistent with the estimates of the equity-only VAR system. Panel A shows that PE forecasts the excess returns on both the value and growth portfolios negatively, with similar coefficients in both equations. The nominal short rate also forecasts the excess return on both portfolios negatively. Table 4 also considers the real return on Treasury bills (or ex-post real short interest rate) and the excess return on long-term bonds. It shows the well known result that the yield spread forecasts bond excess returns positively, with a t-statistic above 4. It also shows that the nominal short rate weakly forecasts bond excess returns with a positive sign. The yield spread and the nominal short rate all forecasts the real return on T-bills. In both tables, the R^2 's of the return forecasting equations are below 5%. This magnitude is

¹⁸Unconstrained VAR estimates and the associated portfolio allocations are available from the authors upon request.

typical of return forecasting equations at monthly frequencies, and it implies considerably larger R^2 's at lower frequencies (Campbell and Thompson, 2005).

Panel B in Table 4 shows the covariance structure of innovations in the equities-and-bonds system. The upper left block of the table shows the covariance structure of unexpected excess returns. Unexpected excess returns on value stocks are highly positively correlated with unexpected excess returns on growth stocks. Both portfolios show similar low positive correlations with the real return on T-bills and bonds. The upper right block shows the correlation structure of unexpected returns with the state variables. We have already discussed the large positive correlation of unexpected returns on growth stocks with PE. Unexpected returns on value stocks are also positively correlated with shocks to PE, but their correlation is smaller than the correlation of growth stocks with shocks to PE. This is consistent with the negative correlation between shocks to PE and the unexpected return on V-G shown in Table 3. We have already noted that this lower correlation implies that value stocks are less correlated with innovations in aggregate stock market discount rates than growth stocks. As we show below, this has important implications for our asset allocation results. Finally, unexpected excess returns on bonds are highly negatively correlated with the short nominal interest rate (Campbell, Chan and Viceira 2003, Campbell and Viceira 2005).

5.2 Portfolio allocations in the equity-only case

We now examine the optimal portfolio allocation of a K -period, power-utility investor who is constrained to invest only in two portfolios of equities—a value stock portfolio and a growth stock portfolio—and behaves as if investment opportunities evolved according to the equity-only VAR of Table 3. Section 4 shows that the optimal portfolio rule for this investor is an affine function of the vector of state variables, with coefficients that change with the investor's investment horizon. To facilitate interpretation, we omit reporting these coefficients and instead present our results in the form of plots with mean percentage portfolio allocations to each asset across different investment horizons. Given the affine form of the portfolio rule, the mean portfolio allocation at any horizon is simply the optimal portfolio rule evaluated at the unconditional mean of the vector of state variables. We consider horizons between 1 month, at the left end of the plots, and 300 months (or 25 years) at the other end.

The 1-month mean percentage allocations correspond to the one-period myopic allocation described in equation (18). We have shown that the myopic portfolio allocation is a weighted average of the one-period mean-variance efficient tangency portfolio (20) and the one-period mean-variance efficient GMV portfolio (21), with weights $1/\gamma$ and $(1 - 1/\gamma)$.

Thus a short-term log utility investor holds only the tangency portfolio, and a short-term, infinitely risk averse investor holds only the GMV portfolio.

Table 5 shows the composition of the tangency portfolio and the GMV portfolio in the equity-only case and in the equities-and-bonds case. The tangency portfolio is heavily tilted toward value stocks. A short-term log utility investor holding this portfolio invests on average about 572% of her wealth in the value portfolio, and finances this long position with a 472% short position in growth stocks. Table 2 and Table 4 help in understanding the value tilt in the tangency portfolio. Table 2 shows that the mean return on the value portfolio is larger than the mean return on the growth portfolio, and Table 4 shows that the returns on the value portfolio are highly positively correlated with the returns on the growth portfolio. This large positive correlation (90%) implies that, from the perspective of a short-term investor, value and growth stocks are close substitutes. Thus a modestly risk averse myopic investor will try to take advantage of the mean return spread between value and growth by taking aggressive long and short positions in value and growth stocks, respectively.

The GMV portfolio is also heavily tilted toward value stocks—about 90% long value stocks. With a short-run return correlation of 90% between the return on the value stock portfolio and the growth stock portfolio, the GMV portfolio is heavily tilted toward the portfolio with the smallest return volatility, which Table 4 shows is the value stock portfolio.¹⁹

Figure 3 presents plots of the mean percentage allocation to growth and value for an investor with coefficients of relative risk aversion of 2 (Panel A), 5, 10 and 20 (Panel B), and 500 (Panel C). These figures show that the composition of the optimal portfolio changes dramatically across investment horizons. The left end of the plots shows the 1-month myopic portfolio allocations which, consistent with the results shown in Table 5, is largely invested in the value stock portfolio with short positions in the growth stock portfolio. But as we consider longer investment horizons, the mean optimal allocation to value decreases—and correspondingly the mean allocation to growth increases—and eventually flattens out.²⁰ Interestingly, the reduction in the value tilt of the portfolio becomes more pronounced as we

¹⁹ Another way to understand this result is to apply equation (21), which describes the GMV portfolio, to the only-equities case. Equation (21) implies that the GMV portfolio allocation to value stocks is equal to the negative of the correlation between the unexpected return on growth stocks with the unexpected return on V-G, times the ratio of return volatilities. Table 3 shows that this correlation is about -41% , and the ratio of volatilities is about 2.3. The product of these two quantities is .93, which is the GMV allocation to value stocks.

²⁰ The pattern is monotonically decreasing for most coefficients of relative risk aversion. However, for some investors there is a slight increase in the value allocation as they move from a 1-month horizon to a 2-month horizon. As we have already noted, this short-term effect is caused by lagged return predictability.

consider increasingly risk averse investors. This reflects an intertemporal hedging demand for growth stocks which is increasing in risk aversion and investment horizon.

Panel C in Figure 3 makes this point clear. This panel plots the mean portfolio allocation of an investor who is effectively infinitely risk averse.²¹ This investor chooses the least risky portfolio at any given investment horizon, regardless of expected return. Of course, at short horizons this is the GMV portfolio, which Table 5 and the leftmost point in the plots show is 90% invested in the value portfolio, and only 10% invested in the growth portfolio. But intertemporal hedging considerations lead highly conservative investors with longer investment horizons to move away from value into growth. For investors with horizons of about 4 years (or 48 months), the least risky portfolio is already 50% in value stocks and 50% in growth stocks. At horizons of about 12 years (or 144 months), the long-run least risky portfolio is about 30% in the value portfolio, and 60% in the growth portfolio. Interestingly, this portfolio is very close to the aggregate stock market portfolio, suggesting that the least risky portfolio of a long-term, all-equity investor has no value tilt or a very small tilt.

In order to best understand the increasing allocation to growth as a function of the agent's risk aversion and investment horizon it is useful to consider the two components of total portfolio demand. Since equation (18) implies that the mean myopic portfolio demand is constant across all investment horizons, the entire variation across investment horizons in total portfolio demand we observe in Figure 3 must be ascribed to the hedging component of portfolio demand. Thus the increasing allocation to growth as investment horizon and risk aversion increase reflects intertemporal hedging considerations. Growth stocks appear to be less risky than value stocks to long-term investors because they help investors hedge temporary fluctuations in aggregate stock market discount rates.

5.3 Portfolio allocations in the presence of bonds and cash

We now examine the optimal portfolio strategies of investors who can invest in long-term nominal bonds and Treasury bills, in addition to the value and the growth stock portfolios. This section assumes that investors view investment opportunities as reflected in the VAR estimates of Table 4. Figures 4, 5 and 6 show the mean portfolio allocations to each asset of investors with different coefficients of relative risk aversion and investment horizons. (The layout of these figures is identical to the asset allocation figures described in Section 5.2.)

Panel A in Figure 4 plots the optimal mean portfolio allocations across investment

²¹We show in the Appendix that the optimal portfolio policy converges to an envelope as $\gamma \rightarrow \infty$. In this particular application, the convergence occurs for values of γ in the vicinity of 500.

horizons of an investor with a coefficient of relative risk aversion equal to 2. This fairly aggressive investor takes long levered positions in the value stock portfolio and the long-term bond, financed with short positions in the growth stock portfolio and cash. In contrast with Figure 3, Panel A in Figure 4 shows that the optimal portfolio allocation to value of aggressive investors does increase with investment horizon when cash and long-term bonds are also available for investing. Interestingly, these investors also choose to increase their exposure to growth stocks as their investment horizon lengthens.

Thus Panel A in Figure 4 indicates that the optimal overall exposure to equities of aggressive long-term investors increases with horizon, consistent with previous findings in the literature on optimal portfolio choice (see, for example, Campbell and Viceira 1999, 2002, 2003). However, it is not obvious from this figure whether it is also optimal for these investors to change the value tilt of their portfolios in any particular direction.

To examine this issue, Panel B in Figure 4 plots the allocation to the overall equity market and the value tilt with respect to the market portfolio implicit in the value and growth allocations. Assuming a 70%-30% split of the market portfolio between growth and value, we compute the mean total stock market exposure of a portfolio at any horizon as $\bar{\alpha}_G/0.7$, where $\bar{\alpha}_G$ is the mean portfolio allocation to the growth stock portfolio, and the mean value tilt of the portfolio as $(0.3/0.7)\bar{\alpha}_G + \bar{\alpha}_V$, where $\bar{\alpha}_V$ is the mean portfolio allocation to the value stock portfolio.²²

Panel B in Figure 4 shows that the portfolios of short-horizon investors have a large value tilt, and an implicit large short position in the market portfolio. As the investment horizon increases, the short position in the market portfolio becomes smaller, but the value tilt remains fairly constant. Thus aggressive long-term investors hedge time variation in investment opportunities by increasing their exposure to the market portfolio, but not at the expense of decreasing their net exposure to value stocks.

Figure 5 and Figure 6 confirm these findings for increasingly conservative investors. For investors willing to hold risky assets in their portfolios, the exposure to the aggregate stock market portfolio increases with the investment horizon, but the value tilt stays constant across investment horizons. Both the overall market exposure and the value tilt of the optimal portfolios decrease as risk aversion increases. Highly risk averse investors such as the investor represented in Figure 6 choose portfolios with no exposure to equities at all horizons. At short horizons, they are fully invested in cash, while at long horizons they are

²²This is consistent with the shares shown in Figure 1. We have also re-estimated the VAR model in Table 3 including the return on the value portfolio and the market portfolio instead of the growth portfolio, and computed optimal allocations to the value portfolio and the market portfolio. The results are almost identical to the results shown in Figure 3. This is not entirely unexpected in light of Figure 1.

fully invested in long-term bonds.

There are three main lessons one can draw from these figures. First, investors with low or moderate levels of risk aversion increase their allocation to equities at long horizons. This increased exposure to equities reflects a positive intertemporal hedging demand for stocks. Campbell and Viceira (1999, 2002, 2003) and others have shown that these investors use the aggregate stock market portfolio to hedge time variation in expected stock returns, since realized stock returns tend to increase at times when expected stock returns decrease.

Second, in a world where investors can invest in growth stocks and value stocks separately, long-term investors hedge time variation in expected aggregate stock returns by increasing their holdings of growth stocks relative to their holdings of value stocks. This allows them to preserve the value tilt of their portfolios at the same time, thus taking advantage of the average return spread between value stocks and growth stocks.

Third, growth stocks play no role in the optimal portfolios of highly risk averse investors when bonds and cash are available for investment, regardless of horizon. Brennan and Xia (2000), Campbell and Viceira (2001), Wachter (2003) and others have shown that highly risk averse investors are interested in hedging time variation in real interest rates, rather than time variation in expected stock returns. For these investors the optimal portfolio strategy is to be fully invested in a zero-coupon, inflation-indexed bond that matches their investment horizon. When this bond is not directly available to them, investors will try to synthesize it out of the menu of available assets.

Figure 6 shows that the portfolio that best synthesizes this bond out of cash, a constant maturity nominal long-term bond, and value and growth equities, is a mix of cash and the nominal long-term bond. The weight of cash in this mix decreases as horizon increases. Growth stocks do not play any role in the portfolio of these investors, suggesting that they are not particularly good hedges of bond returns.

5.4 The value of dynamic rebalancing

Our empirical analysis so far has focused on the portfolio allocations to value and growth stocks implied by the optimal dynamic portfolio strategy (26) of Section 4. Under this strategy, investors dynamically vary their portfolio shares each period in response to changes in investment opportunities. They also change portfolio shares as their investment horizon shortens, even if investment opportunities do not change from one period to the next—this is the effect captured by the $A_i^{(\tau)}$ coefficients and reflected in the mean allocation plots in Section 5.

It is useful to compare this strategy to the simpler constrained strategy (15) of Section 3, where investors optimally choose portfolio weights at the beginning of their investment horizon and rebalance to these weights thereafter. Since this strategy does not allow the investor to modify her allocation across time beyond rebalancing back to the vector of fixed portfolio weights, it must provide an optimal one-shot response to the term-structure of the risk-return tradeoff at the relevant investment horizon. We have argued in Section 3 that the constrained constant proportion strategy provides an intuitive and practically relevant benchmark to analyze the benefits of optimal dynamic rebalancing to investors.

Figure 7 plots the optimal portfolio allocations to value and growth implied by the dynamic rebalancing strategy (solid line) and the constant proportion strategy (dash-dot line) of all-equity investors with an investment horizon of 300 months (or 25 years) and coefficients of relative risk aversion equal to 2 (left panel) and 500 (right panel). The figure plots the allocations along a 300 period path, assuming that the state vector stays at its unconditional mean.

Figure 7 shows that unconstrained long-horizon aggressive investors hold on average more value stocks in their portfolios than their constrained counterparts, while long-horizon conservative investors choose initial portfolios which are less tilted toward value than their constrained counterparts. We show below that aggressive investors value very highly the ability to time equity markets, and they are more willing to take on value risk when they know that they will be able to change their future value exposure in response to adverse changes in investment opportunities. By contrast, highly conservative long-term investors change their portfolios very little in response to changes in expected returns. These investors choose to move their portfolios toward value and away from growth over time as their investment horizon shortens, regardless of how expected returns change along the way. When constrained to hold constant portfolio weights over time, they choose weights which average out this horizon effect.

Of course, since the constant-proportion strategy is constrained relative to the dynamic portfolio sequence, it must always lead to a welfare loss, even if it appears to deliver superior returns to the optimal dynamic strategy over some sub-sections of the investment plan. We now examine this welfare loss in detail. To do so we take advantage of the fact that the VAR(1) class of models is closed under temporal aggregation—see proof in the Appendix. This allows us to aggregate the systems we estimate using monthly data to any desired frequency, and to examine the value of the rebalancing frequency to the investor, for which the constant-proportion strategy is the relevant limiting case.

Table 6 explores the effect of decreasing the rebalancing frequency on the welfare of investors with different investment horizons and risk aversion coefficients in the context

of the equity-only case. Our base case is the dynamic strategy (26) with a monthly rebalancing periodicity, which we extend gradually to a quarter, half-year, and one year. Finally, we also consider the constrained constant proportion strategy, where the investor rebalances each period, but only to fixed weights optimally chosen at the beginning of her horizon. Panel A shows the certainty equivalent of the dynamically rebalanced strategy (15) in the equity-only case. Panel B reports the welfare losses of dynamic strategies with lower rebalancing frequencies, and the welfare loss of the constant-proportion strategy. All welfare losses are computed as the percentage difference between the certainty equivalent of the limited rebalancing strategy and the certainty equivalent of the dynamically rebalanced strategy.²³ Panel C reports the largest monthly management fee which an investor constrained to follow a strategy with limited rebalancing would pay to gain access to the monthly dynamically rebalanced strategy. These fees re-scale the welfare losses in Panel B to facilitate comparisons across investment horizons.²⁴

Panel B shows that the welfare loss from not rebalancing dynamically at a monthly frequency increases as we lower the rebalancing frequency. At long horizons, welfare losses are maximal for investors who follow the constant proportion strategy. Controlling for rebalancing frequency, total welfare loss rises as the investment horizon increases (more time to exploit return predictability) and as risk-aversion decreases (more aggressive portfolio positions). At long horizons, welfare losses are large for all investors, but they are particularly large for aggressive investors. At horizons of 300 months (or 25 years), the welfare loss is 61.9% for log utility investors and a fourth of that (a large 15.8%) for investors with coefficients of relative risk aversion of 10.

The normalization of welfare losses shown in Panel C suggests however that welfare losses are roughly similar across investment horizons for investors who are able to rebalance their portfolios dynamically. For investors constrained to follow the constant propor-

²³In both strategies we hold weights constant between rebalancing dates, and we use simulated data at a monthly frequency. To compute the certainty equivalent measure for each rebalancing frequency we simulate 20,000 paths of the VAR system shown in Table 2. Next we compute the optimal strategy (26) or (15) along each path. This allows us to compute the realized terminal wealth and realized utility of terminal wealth for each path. We then average realized utility across paths to obtain our measure of expected utility. This measure is subsequently inverted to obtain a certainty equivalent of wealth. We use the same set of simulated paths to evaluate combinations of the investment horizon, risk aversion and rebalancing frequency, and for each rebalancing frequency we appropriately perform a temporal aggregation of the VAR system.

²⁴Specifically, the fee ϕ is computed as the solution to

$$(1 + g)(1 - \phi) = (1 + g_{const}),$$

where g and g_{const} are the monthly growth rates implied by the certainty equivalent of wealth of the monthly dynamic rebalancing strategy and the constrained rebalancing strategy respectively. If W_{CE} denotes the certainty equivalent of wealth, the growth rate is equal to $(W_{CE})^{1/K} - 1$.

tions strategy, management fees increase with investment horizon.²⁵ Panel C shows that a constant-proportion log utility investor would be willing to pay a monthly fee equal to 0.27% of her wealth per month (3.24% per year) to be able to dynamically rebalance her portfolio monthly. A constant-proportion investor with a coefficient of risk aversion of 10 would be willing to pay a much smaller monthly fee of 0.06% (or about 0.7% per year) to get access to the monthly dynamic rebalancing strategy.

In general, Table 6 shows that long-horizon, log utility investors are the investors who suffer the largest welfare losses from not rebalancing dynamically. These investors change their portfolios in response to changes in expected returns, but choose their portfolios without any intertemporal hedging considerations. Thus for these investors welfare losses are all about the timing value of predictability. For more risk averse investors, intertemporal hedging considerations also influence their optimal portfolio decisions. The relatively small welfare losses for these investors from not rebalancing their portfolios dynamically suggest that intertemporal hedging portfolio demands are relatively stable over time and change only slowly in response to changes in expected returns. Table 7 illustrates this intuition.

Table 7 reports the monthly standard deviation of the portfolio allocation to value—the volatility of the portfolio allocation to growth is identical²⁶—and the ratio of the standard deviation of the intertemporal hedging demand to value over the standard deviation of the total portfolio demand for investors with coefficients of relative risk aversion of 5, 10, and 500, and investment horizons of 12 months, 60 months, 120 months, and 300 months. This table shows that the volatility of total portfolio demand does not change across investment horizons, but it declines dramatically as we consider increasingly risk averse investors. For example, the monthly volatility of the portfolio of an investor with a coefficient of relative risk aversion of 2 is about 400% per month, while the monthly volatility of an investor with a coefficient of relative risk aversion of 500 is less than 2%. Most interestingly, this table shows that most of portfolio volatility is caused by the variability of the myopic component. The fraction of total portfolio volatility due to intertemporal hedging is very low for all investors regardless of their risk aversion: It is always smaller than 5%. Thus while intertemporal hedging accounts for all the variability of portfolio demand across investment horizons, myopic portfolio demand accounts for almost all of the variability of portfolio demand over time. This helps understand why welfare losses from following the constrained constant proportion strategy (15) or a low frequency dynamic rebalancing policy are much smaller

²⁵This is not surprising though in light of the fact that a dynamic rebalancing strategy is time consistent—i.e., we can interpret the numbers in each row of the panel as describing the same investor as her horizon changes—while the constant proportion strategy is not—i.e., each number in each row of Panel B.4 or C.4 describes a different investor.

²⁶To see this note that the myopic portfolio demands add up to one, and the intertemporal hedging demands add up to zero. Thus the total portfolio allocation to growth is just one minus the total portfolio allocation to value.

for highly conservative all-equity investors than for aggressive all-equity investors.

Next we turn to the welfare analysis of the case when the investable universe is expanded to include T-bills and the 5-year bond. This is shown in Table 8, whose structure is identical to Table 6. Although we find an identical set of patterns of welfare losses across rebalancing frequencies, investment horizons and levels of risk aversion, their magnitudes are much larger than in the all-equities case. For example, for investors with a 25-year investment horizon the welfare loss from following the constant proportions strategy ranges from 99.9% (when $\gamma = 1$) to 73.2% (when $\gamma = 10$). These differences in the certainty equivalents of wealth imply maximal management monthly fees of 2.44% and 0.44% per month, respectively. That is, a log-utility constant proportion agent would be willing to pay a management fee of up to 29.3% per year to gain access to the optimal dynamic strategy, and the investor with coefficient of relative risk aversion of 10 would be willing to pay 5.3% per year. This result suggests that with an expanded investment set the constant-proportion strategy represents a very poor approximation to the optimal dynamic portfolio rule. A comparison of the certainty equivalent of wealth shown in Panel A of Table 6 with the one shown in Panel A of Table 8 also shows that the ability to access cash and bonds is extremely valuable to all investors.

5.5 Conclusion

This paper makes several contributions to the theory and practical implementation of modern portfolio theory. First, it examines the portfolio choice problem of a long-term investor with isoelastic utility of wealth at a finite horizon, who faces time variation in expected returns and interest rates described by a first-order vector autoregressive system. The paper develops closed form solutions for two cases of practical and theoretical interest. The first case is a constant proportion strategy in which investors optimally choose their investment horizon portfolio weights at the beginning of the investment period, and rebalance their portfolio back to those weights at each point in time. This strategy captures the spirit of the investment policies often adopted by institutional and individual investors. The second case is an unconstrained strategy where investors rebalance their portfolios each period, but not necessarily back to fixed weights. We show that the optimal dynamic rebalancing strategy is affine in the vector of state variables describing investment opportunities, with coefficients that change with investment horizon. This solution is based on an approximation to the log return on wealth which becomes increasingly accurate as the frequency of rebalancing increases, and it is exact in continuous time. An important advantage of this solution is that it can easily and quickly evaluate investment opportunity sets with any number of assets and state variables. By contrast, existing analytical models do not allow for simultaneous time variation in risk premia and interest rates—with the exception of the

approximate solution in Campbell, Chan and Viceira, (2003)—, and traditional numerical solution methods, which can handle both, become difficult to apply in problems with more than a few state variables. Thus our solution methods provides a step forward toward the practical implementation of dynamic portfolio choice models with realistically complex investment opportunity sets.

Second, we apply our methods to study the practically relevant problem of value and growth investing at long horizons. We find that on average equity-only investors with short horizons optimally choose portfolios heavily tilted toward value and away from growth, regardless of their risk aversion. Aggressive short-term investors find it optimal to hold long large positions in value stocks offset by large short positions in growth stocks because the mean return spread between value and growth is positive, and their returns are highly positively correlated. Highly risk averse short-term investors hold large positions in value stocks because of their smaller return volatility and high correlation with growth. However, the optimal allocation to value decreases dramatically—and correspondingly the optimal allocation to growth increases—for investors with longer horizons. This effect is strongest for long-horizon, highly risk averse investors, who hold large long positions in growth stocks. The increasing portfolio demand for growth stocks across investment horizons is driven by intertemporal hedging motives. Growth stocks are better suited than value stocks to hedge against adverse changes in investment opportunities in the equity market, because they are more highly negatively correlated with changes in aggregate stock discount rates than value stocks are. Thus long-horizon investors find value stocks riskier than growth stocks, and see the unconditional value spread as a risk premium for bearing this risk.

We also consider investors who can choose among value stocks, growth stocks, cash, and long-term bonds. In this case, we find that investors willing to hold equities in their portfolios still increase their allocation to growth stocks as their investment horizon lengthens. This demand is once again driven by intertemporal hedging motives. However, in this case they have the ability to increase their portfolio demand for growth stocks without sacrificing their exposure to value stocks. We find that the value tilt—relative to the aggregate market exposure—of these investors remains fairly stable across investment horizons. Naturally, as we consider increasingly risk averse investors, they shift out of equities and into fixed income instruments. Short-horizon, highly risk averse investors hold large positions in cash, while long-horizon highly risk averse investors hold large positions in long-term bonds.

Third, we compare the welfare loss to investors from using constrained portfolio rules instead of the unconstrained dynamic rebalancing strategy. We consider the welfare loss from a constant proportion strategy relative to the optimal dynamic strategy as we vary the horizon and rebalancing frequencies. For all-equity investors, we find that the welfare losses from the constant proportion strategy and from the strategy of infrequent rebalancing

are significant for long-horizon investors with low risk aversion coefficients. By contrast, moderately risk averse, long horizon all-equity investors do not face large welfare losses from adopting constant proportion or infrequent rebalancing strategies that take into account their investment horizon and the effect of return predictability and changing interest rates on risk. For long-horizon investors who can invest in equities, bonds, and cash, welfare losses from adopting investment policies with fixed portfolio weights or portfolio weights that move over time infrequently are large, regardless of their risk aversion coefficient.

This paper develops a flexible analytical framework for the study of dynamic portfolio choice problems with a large number of assets and state variables. This flexibility is enhanced by modeling investment opportunities as a VAR(1), which is extremely easy to estimate from data. However, the model has important limitations that should be kept in mind. First, the model ignores intermediate consumption or labor income, which are relevant for individual investors. Campbell, Chan, and Viceira (2003) also study a high-dimensional portfolio choice problem with intermediate consumption, but their solution is accurate only when the optimal consumption-wealth ratio is not too variable. Second, our model does not consider short-sales or borrowing constraints. Allowing for either of these constraints would take us outside the realm of analytically tractable models, and would require a fully numerical solution which might be difficult to implement in high-dimensional problems. Constraints typically ameliorate the extreme positions that unconstrained models sometimes recommend based on the capital market assumptions of the investor, but they do not often change the direction of those recommendations. Thus an unconstrained solution is an essential step in any asset allocation problem, since it allows for an examination of the economic implications of the investor's view of investment opportunities. Third, when the VAR parameters are estimated from data, we assume that investors take them at face value, ignoring estimation uncertainty or learning (Brennan 1998, Xia 2001). Of course, one could consider investors who weight the empirical likelihood function with priors that are not fully informative. Previous research suggests that, as long as investor's priors are not completely diffuse, allowing for parameter uncertainty typically results in portfolio allocations which are a smoothed version of the allocations that result from ignoring it (Kandel and Stambaugh 1996, Barberis 2000, Wachter and Warusawitharana 2005). Investigation of highly dimensional models with labor earnings uncertainty, portfolio constraints, and parameter uncertainty and learning are important directions for future research.

6 References

Balduzzi, Perluigi and Anthony Lynch, 1999, "Transaction Costs and Predictability: Some Utility Cost Calculations", *Journal of Financial Economics* 52, 47–78.

- Barberis, Nicholas C., 2000, "Investing for the Long Run When Returns Are Predictable", *Journal of Finance* 55, 225–264.
- Brandt, Michael, 1999, "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach", *Journal of Finance* 54, 1609–1645.
- Brennan, Michael J., 1998, "The Role of Learning in Dynamic Portfolio Decisions", *European Finance Review* 1, 295–306.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, "Strategic Asset Allocation", *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1999, "The Use of Treasury Bill Futures in Strategic Asset Allocation Programs", in William T. Ziemba and John M. Mulvey eds. *World Wide Asset and Liability Modeling*, Cambridge University Press, 205–228.
- Brennan, Michael J., Ashley Wang and Yihong Xia, 2004, "Estimation and Test of a Simple Model of Intertemporal Asset Pricing," *Journal of Finance* 59, 1743-1775.
- Brennan, Michael J. and Yihong Xia, 2000, "Stochastic Interest Rates and the Bond-Stock Mix," *European Finance Review* 2000, 4, 197-210.
- Brennan, Michael J. and Yihong Xia, 2001, "Assessing Asset Pricing Anomalies", *Review of Financial Studies* 14, 905-942.
- Brennan, Michael J. and Yihong Xia, 2002, "Dynamic Asset Allocation Under Inflation", forthcoming *Journal of Finance*.
- Campbell, John Y., 1987, "Stock Returns and the Term Structure", *Journal of Financial Economics* 18, 373–399.
- Campbell, John Y. and Robert J. Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors", *Review of Financial Studies* 1, 195–228.
- Campbell, John Y. and Robert J. Shiller, 1998, "Valuation Ratios and the Long-Run Stock Market Outlook", with Robert J. Shiller, *Journal of Portfolio Management* 24(2), 11–26.
- Campbell, John Y. and Robert J. Shiller, 2005, "Valuation Ratios and the Long-Run Stock Market Outlook: An Update," NBER Working Paper No. 8221, forthcoming in Nicholas Barberis and Richard Thaler eds. *Advances in Behavioral Finance*, Volume II, Princeton University Press.

- Campbell, John Y. and Robert J. Shiller, 1991, “Yield Spreads and Interest Rates: A Bird’s Eye View”, *Review of Economic Studies* 58, 495–514.
- Campbell, John Y. and Robert J. Shiller, 1996, “A Scorecard for Indexed Government Debt”, in Ben S. Bernanke and Julio Rotemberg eds. *National Bureau of Economic Research Macroeconomics Annual* 11:155–197.
- Campbell, John Y. and Thompson, 2005, “Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average?,” NBER Working Paper No. 11468.
- 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, 2000, “Consumption and Portfolio Decisions when Expected Returns are Time Varying: Erratum”, unpublished paper, Harvard University, available on the authors’ websites.
- Campbell, John Y. and Luis M. Viceira, 2001, “Who Should Buy Long-Term Bonds?”, *American Economic Review* 91, 99–127.
- Campbell, John Y. and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, Oxford.
- Campbell, John Y. and Luis M. Viceira, 2005, “The Term Structure of the Risk-Return Tradeoff,” *Financial Analysts Journal* 61, 34-44.
- Campbell, John Y. and Tuomo Vuolteenaho, 2004, “Bad Beta, Good Beta”, *American Economic Review* 94, 1249-1275.
- Campbell, John Y. and Motohiro Yogo, 2005, “Efficient Tests of Stock Return Predictability,” forthcoming *Journal of Financial Economics*.
- Canner, Niko, N. Gregory Mankiw, and David N. Weil, 1997, “An Asset Allocation Puzzle”, *American Economic Review* 87, 181–191.
- Chacko, George and Luis M. Viceira, 2005, “Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets”, forthcoming *Review of Financial Studies*.
- Chan, Y. Lewis, John Y. Campbell, and Luis M. Viceira, 2003, “A Multivariate Model of Strategic Asset Allocation,” *Journal of Financial Economics*, 67, 41-80.
- Elleswarapu and Reinganum, 2004, “The predictability of aggregate stock market returns: Evidence based on glamour stocks,” *Journal of Business*.

- Fama, Eugene F. and Robert R. Bliss, 1987, "The Information in Long-Maturity Forward Rates," *American Economic Review*, 77, 680-92.
- Fama, Eugene and Kenneth R. French, 1988, "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics* 22, 3-27.
- Fama, Eugene and Kenneth R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds", *Journal of Financial Economics* 25, 23-49.
- Fama, Eugene F. and Kenneth R. French, 1992, "The cross-section of expected stock returns," *Journal of Finance* 47, 427-465.
- Fama, Eugene F. and Kenneth R. French, 1993, "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics* 33, 3-56.
- Fama, Eugene and Kenneth R. French, 1996, "Multi-factor Explanations of Asset Pricing Anomalies," *Journal of Finance* 51, 55-84.
- Fama, Eugene and G. William Schwert, 1977, "Asset Returns and Inflation", *Journal of Financial Economics* 5, 115-146.
- Glosten, Lawrence R., Ravi Jagannathan, and David Runkle, 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, William N. and Philippe Jorion, 1993, "Testing the Predictive Power of Dividend Yields", *Journal of Finance* 48, 663-679.
- Hodrick, Robert J., 1992, "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement", *Review of Financial Studies* 5, 357-386.
- Kandel, Shmuel and Robert F. Stambaugh, 1996, "On the Predictability of Stock Returns: An Asset Allocation Perspective", *Journal of Finance* 51, 385-424.
- Kim, Tong Suk and Edward Omberg, 1996, "Dynamic Nonmyopic Portfolio Behavior", *Review of Financial Studies* 9, 141-161.
- Liu, Jun, 1998, "Portfolio Choice in Stochastic Environments", unpublished paper, Stanford University.
- Lynch, Anthony W., 2001, "Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability", *Journal of Financial Economics* 62, 67-130.
- Markowitz, Harry, 1952, "Portfolio Selection", *Journal of Finance* 7, 77-91.

- Merton, Robert C., 1969, “Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case”, *Review of Economics and Statistics* 51, 247–257.
- Merton, Robert C., 1971, “Optimum Consumption and Portfolio Rules in a Continuous-Time Model”, *Journal of Economic Theory* 3, 373–413.
- Merton, Robert C., 1973, “An Intertemporal Capital Asset Pricing Model”, *Econometrica* 41, 867–87.
- Merton, Robert C., 1990, *Continuous Time Finance*, Basil Blackwell, Cambridge, MA.
- Samuelson, Paul A., 1969, “Lifetime Portfolio Selection by Dynamic Stochastic Programming”, *Review of Economics and Statistics* 51, 239–246.
- Shiller, Robert J., John Y. Campbell, and Kermit L. Schoenholtz, 1983, “Forward Rates and Future Policy: Interpreting the Term Structure of Interest Rates”, *Brookings Papers on Economic Activity* 1, 173–217.
- Sorensen, Carsten, and Anders B. Trolle, 2005, “Dynamic Asset Allocation and Latent Variables”, working paper, Copenhagen Business School.
- Wachter, Jessica, 2002, “Portfolio and Consumption Decisions Under Mean-Reverting Returns: An Exact Solution for Complete Markets”, forthcoming *Journal of Financial and Quantitative Analysis*.
- Wachter, Jessica and Warusawitharana, 2005, “Predictable Returns and Asset Allocation: Should a Skeptical Investor Time the Market?,” working paper, Wharton School of Business, University of Pennsylvania.
- Xia, Yihong, 2001, “Learning About Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation”, *Journal of Finance* 56, 205–246.

A Appendix: Model setup

In this section we introduce the VAR(1) specification describing the evolution of the risky asset returns and the state variables used to forecast them. We then present the Campbell-Viceira (2003) approximation to the log portfolio return, state the formulas relating the single-period VAR moments to their long-horizon counterparts, and introduce the selector matrix notation used in the ensuing derivations.

A.1 VAR(1) Model

The vector being modeled as a first order auto-regression, \mathbf{z}_{t+1} , is comprised of the log return on the benchmark asset ($r_{1,t+1}$), the excess log returns on the n risky assets computed in relation to the benchmark asset (\mathbf{x}_{t+1}) and the m state variables (\mathbf{s}_{t+1}), which characterize the investment opportunity set and are assumed to have predictive power for the asset returns. The dynamics of the vector auto-regression are given by:

$$\mathbf{z}_{t+1} = \Phi_0 + \Phi_1 \mathbf{z}_t + \mathbf{v}_{t+1} \quad (28)$$

where \mathbf{z}_{t+1} is:

$$\mathbf{z}_{t+1} = \begin{bmatrix} r_{1,t+1} \\ \mathbf{x}_{t+1} \\ \mathbf{s}_{t+1} \end{bmatrix} \quad (29)$$

The VAR(1) model is parameterized by: an $(1+n+m) \times 1$ vector of intercepts - Φ_0 , and an $(1+n+m) \times (1+n+m)$ matrix of slope coefficients - Φ_1 . The $(1+n+m) \times 1$ vector of shocks \mathbf{v}_{t+1} has a mean zero multivariate normal distribution with covariance matrix:

$$\Sigma_v = \begin{bmatrix} \sigma_1^2 & \sigma'_{1x} & \sigma'_{1s} \\ \sigma_{1x} & \Sigma_{xx} & \Sigma'_{xs} \\ \sigma_{1s} & \Sigma_{xs} & \Sigma_{ss} \end{bmatrix} \quad (30)$$

A.2 Portfolio return approximation

Following Campbell and Viceira (QJE 1999), we approximate the log return on the portfolio as:

$$r_{p,t+1} = \alpha'_t \mathbf{x}_{t+1} + r_{1,t+1} + \frac{1}{2} \alpha'_t (\sigma_x^2 - \Sigma_{xx} \alpha_t) \quad (31)$$

where $\sigma_x^2 \equiv \text{diag}(\Sigma_{xx})$, is the vector consisting of the diagonal elements of Σ_{xx} , the variances of the excess returns. The accuracy of this approximation increases with the rebalancing frequency, and is exact in the limit of continuous trading.

A.3 Conditional K-period moments

In order to assess the value an agent places on the ability to re-balance his portfolio at intermediate dates, when his objective is to maximize the utility of terminal wealth, it is necessary to consider some benchmark portfolio strategies where the opportunity for re-balancing is restricted. The two most intuitive cases are:

one period (myopic) portfolio choice based on K-period return moments and a strategy in which the portfolio weights are chosen at inception and held constant across time (i.e. re-balanced each period to fixed proportions). Because we examine portfolio choice decisions over relatively long horizons the latter strategy is more plausible, as agents are unlikely to allow their weights to drift significantly without intervening. Furthermore, our analytical solutions are based on an approximation to the log portfolio return, which is only accurate at short horizons.

Both of the benchmark strategies require the computation of the conditional K-period moments of the VAR variables, and we present these results here for completeness:²⁷

$$E_t \left[\sum_{i=1}^K \mathbf{z}_{t+i} \right] = \left[\sum_{i=0}^{K-1} (K-i) \Phi_1^i \right] \Phi_0 + \left[\sum_{j=1}^K \Phi_1^j \right] \mathbf{z}_t \quad (32)$$

$$Var_t \left[\sum_{i=1}^K \mathbf{z}_{t+i} \right] = Var_t \left[\sum_{q=1}^K \left[\sum_{p=0}^{K-q} \Phi_1^p v_{t+q} \right] \right] \quad (33)$$

In the important case of no predictability ($\Phi_1 = \mathbf{0}$) it is trivial to verify that the moment equations simplify to:

$$\frac{1}{K} E_t \left[\sum_{i=1}^K \mathbf{z}_{t+i} \right] = \Phi_0 \quad \frac{1}{K} Var_t \left[\sum_{i=1}^K \mathbf{z}_{t+i} \right] = \Sigma_v \quad (34)$$

A.4 Selector matrix notation

In order to facilitate matrix operations in the remainder of the derivations we introduce the *selector matrix notation*. A selector matrix is simply a matrix that, when post-multiplied by the target matrix, returns a set of rows and columns from the target. A selector matrix (operator), \mathbf{H} , has equal row and column size as the target matrix and is comprised of zeros and ones. An operator with a single subscript, e.g. \mathbf{H}_i , selects the rows corresponding to the set of variables i from the target matrix; an operator with a double subscript \mathbf{H}_{ij} first selects the rows corresponding to variable i and then the columns corresponding to variable j . In other words if we denote the target matrix \mathbf{T} , the double selection operation can be expressed as:

$$\mathbf{H}_{ij} \mathbf{T} = (\mathbf{H}_j (\mathbf{H}_i (\mathbf{T})))'$$

Further, we denote the result of selecting the rows and columns corresponding to the set of variables ij from the target matrix, \mathbf{T} , as follows:

$$\mathbf{H}_{ij} \mathbf{T} = \mathbf{T}_{ij}$$

If the target matrix already has a subscript (e.g. Φ_0 or Φ_1) the selected rows and columns will be indicated with a superscript.

²⁷A detailed derivation can be found in Campbell and Viceira (2005).

B Appendix: The constant proportion strategy

In order to provide a realistic benchmark for the evaluation of the optimal dynamic portfolio rebalancing strategy we propose a strategy in which the investor is constrained to rebalance her portfolio each period to a fixed vector of weights. The fixed vector of weights is chosen optimally at time t and is applied until the investment horizon at time $t + K$. We refer to this portfolio policy as the *constant proportion* strategy, and denote the optimal weight vector by $\bar{\alpha}_t^{(K)}$. The subscript denotes the time at which the portfolio weight vector is chosen, and the superscript the length of the investment period (i.e. the time to the horizon). Since re-balancing occurs after each period the Campbell-Viceira approximation to the log portfolio return is likely to be very accurate, facilitating a closed-form derivation of the optimal portfolio choice strategy.

The agent's objective is to maximize the time t expected utility of wealth at horizon $t+K$, $E_t[U(W_{t+K}(\boldsymbol{\alpha}))]$. We begin first by writing the terminal wealth as a function of the portfolio weight vector, $\bar{\alpha}_t^{(K)}$:

$$\begin{aligned} W_{t+K}(\bar{\alpha}_t^{(K)}) &= W_t \cdot \exp\left(\sum_{i=1}^K r_{p,t+i}\right) = W_t \cdot \exp\left(\sum_{i=1}^K \left(\bar{\alpha}_t^{(K)'} \mathbf{x}_{t+i} + r_{1,t+i} + \frac{1}{2} \bar{\alpha}_t^{(K)'} \left(\boldsymbol{\sigma}_x^2 - \boldsymbol{\Sigma}_{xx} \bar{\alpha}_t^{(K)}\right)\right)\right) \\ &= W_t \cdot \exp\left(\bar{\alpha}_t^{(K)'} \mathbf{x}_{t \rightarrow t+K} + r_{1,t \rightarrow t+K} + \frac{K}{2} \bar{\alpha}_t^{(K)'} \left(\boldsymbol{\sigma}_x^2 - \boldsymbol{\Sigma}_{xx} \bar{\alpha}_t^{(K)}\right)\right) \end{aligned}$$

where we have introduced the following notation for the K -period log excess returns, and the K -period log return on the reference asset, $\mathbf{x}_{t \rightarrow t+K} = \sum_{i=1}^K \mathbf{x}_{t+i}$ and $r_{1,t \rightarrow t+K} = \sum_{i=1}^K r_{1,t+i}$. The explicit maximization therefore takes the form:

$$\max_{\bar{\alpha}_t^{(K)}} E_t \left[\frac{W_{t+K}^{1-\gamma}}{1-\gamma} \right] = \max_{\bar{\alpha}_t^{(K)}} \frac{W_t^{1-\gamma}}{1-\gamma} \cdot E_t \left[\exp\left(\left(1-\gamma\right) \left(\bar{\alpha}_t^{(K)'} \mathbf{x}_{t \rightarrow t+K} + r_{1,t \rightarrow t+K} + \frac{K}{2} \bar{\alpha}_t^{(K)'} \left(\boldsymbol{\sigma}_x^2 - \boldsymbol{\Sigma}_{xx} \bar{\alpha}_t^{(K)}\right)\right)\right) \right]$$

Since the expression appearing within the expectation operator is distributed lognormally it can be re-written in terms of the first two moments of the its logarithm:

$$\begin{aligned} E_t [(1-\gamma)r_{p,t \rightarrow t+K}] &= (1-\gamma) \cdot \left(\bar{\alpha}_t^{(K)'} E_t[\mathbf{x}_{t \rightarrow t+K}] + E_t[r_{1,t \rightarrow t+K}] + \frac{K}{2} \bar{\alpha}_t^{(K)'} \left(\boldsymbol{\sigma}_x^2 - \boldsymbol{\Sigma}_{xx} \bar{\alpha}_t^{(K)}\right) \right) \\ \text{Var}_t [(1-\gamma)r_{p,t \rightarrow t+K}] &= (1-\gamma)^2 \cdot \left(\bar{\alpha}_t^{(K)'} \boldsymbol{\Sigma}_{xx}^{(K)} \bar{\alpha}_t^{(K)} + \left(\sigma_{1x}^{(K)}\right)^2 + 2 \bar{\alpha}_t^{(K)'} \sigma_{1x}^{(K)} \right) \end{aligned}$$

Making use of these expressions allows us to re-state the maximization problem in its final form:

$$\max_{\bar{\alpha}_t^{(K)}} E_t \left[\frac{W_{t+K}^{1-\gamma}}{1-\gamma} \right] = \max_{\bar{\alpha}_t^{(K)}} \frac{W_t^{1-\gamma}}{1-\gamma} \cdot \left(E_t [(1-\gamma)r_{p,t \rightarrow t+K}] + \frac{1}{2} \text{Var}_t [(1-\gamma)r_{p,t \rightarrow t+K}] \right) \quad (35)$$

Taking derivatives with respect to $\bar{\alpha}_t^{(K)}$ leads to the following first order condition:

$$E_t[\mathbf{x}_{t \rightarrow t+K}] + \frac{K}{2} \boldsymbol{\sigma}_x^2 - K \cdot \boldsymbol{\Sigma}_{xx} \bar{\alpha}_t^{(K)} + (1-\gamma) \left(\boldsymbol{\Sigma}_{xx}^{(K)} \bar{\alpha}_t^{(K)} + \sigma_{1x}^{(K)} \right) = 0$$

with the solution:

$$\bar{\alpha}_t^{(K)} = \left((\gamma-1) \boldsymbol{\Sigma}_{xx}^{(K)} + K \cdot \boldsymbol{\Sigma}_{xx} \right)^{-1} \left(E_t[\mathbf{x}_{t \rightarrow t+K}] + \frac{K}{2} \boldsymbol{\sigma}_x^2 + (1-\gamma) \sigma_{1x}^{(K)} \right) \quad (36)$$

This benchmark solution turns out to have a variety of interesting corner cases, which intersect our more general solution to the dynamic portfolio choice problem. For example, in the case when $K = 1$, the

fixed-proportions strategy truncates to a myopic portfolio choice rule:

$$\bar{\alpha}_t^{(K=1)} = \frac{1}{\gamma} \Sigma_{xx}^{-1} \left(E_t[\mathbf{x}_{t+1}] + \frac{1}{2} \sigma_x^2 + (1 - \gamma) \sigma_{1x} \right) \quad (37)$$

The case of the log-utility investor ($\gamma = 1$) leads to a different simplification:

$$\begin{aligned} \bar{\alpha}_t^{(K)}(\gamma = 1) &= \Sigma_{xx}^{-1} \left(\frac{E_t[\mathbf{x}_{t \rightarrow t+K}]}{K} + \frac{1}{2} \sigma_x^2 \right) \\ &= \Sigma_{xx}^{-1} \left(\frac{1}{K} \mathbf{H}_x \left(\left[\sum_{i=0}^{K-1} (K-i) \Phi_1^i \right] \Phi_0 + \left[\sum_{j=1}^K \Phi_1^j \right] \mathbf{z}_t \right) + \frac{1}{2} \sigma_x^2 \right) \end{aligned} \quad (38)$$

For the fixed-proportions log-utility investor, the portfolio allocation depends on the (per-period) conditional K -period log return, and a volatility correction term. The correction term, however, is different from the one that would appear if the expression in the brackets were to be equivalent to the (per period) conditional K -period arithmetic return. In the latter case the Jensen's correction term would reflect the per-period volatility of the K -period log return, rather than the volatility of a one period log return. The only case in which these two correction terms would be identical is in the case of no return predictability ($\Phi_1 = \mathbf{0}$):

$$\bar{\alpha}_t^{(K)}(\gamma = 1, \Phi_1 = \mathbf{0}) = \Sigma_{xx}^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) \quad (39)$$

since when there is no predictability the volatility terms scale linearly in time, $\sigma_x^2 = \frac{1}{K} (\sigma_x^{(K)})^2$. As is well known, in the absence of return predictability the portfolio choice decision is myopic, and the portfolio weights are invariant to the investment horizon. This conclusion applies more generally to any investor with CRRA preferences. When returns are not predictable, the agent's optimal portfolio satisfies:

$$\bar{\alpha}_t^{(K)}(\Phi_1 = \mathbf{0}) = \frac{1}{\gamma} \Sigma_{xx}^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1 - \gamma) \sigma_{1x} \right) \quad (40)$$

which we show - in the ensuing section - is equivalent to the dynamic portfolio choice strategy with no predictability.

C Appendix: The optimal dynamic strategy

The K-period dynamic portfolio choice problem can be solved, without loss of generality, by a recursive extension of the solution for the optimization problem of the last two periods. In order to arrive at the recursive solution we proceed in two steps. First we derive the portfolio rule in the last period (the base case for policy function recursion) and the associated value function (the base case for the value function recursion). Then we solve the problem in the preceding period, as a function of the value and policy function coefficients from the last period. This enables us to isolate the recursive relationship linking the policy function and value function recursions for two adjacent periods.

For the two period case, the investor's objective remains to maximize:

$$\begin{aligned} E_{t+K-2} \left[\frac{1}{1-\gamma} W_{t+K}^{1-\gamma} \right] &= \frac{1}{1-\gamma} W_{t+K-2}^{1-\gamma} E_{t+K-2} \left[(1 + R_{p,t+K-1})^{1-\gamma} (1 + R_{p,t+K})^{1-\gamma} \right] \\ &= \frac{1}{1-\gamma} W_{t+K-2}^{1-\gamma} E_{t+K-2} \left[\exp \{ (1-\gamma) r_{p,t+K-1} \} E_{t+K-1} \left[\exp \{ (1-\gamma) r_{p,t+K} \} \right] \right] \end{aligned}$$

with respect to the sequence of portfolio choices $\{\alpha_{t+K-2}^{(2)}, \alpha_{t+K-1}^{(1)}\}$. The control variables are indexed by the time at which they are to be applied (subscript) and the time remaining to the horizon (superscript). This convention will allow us to distinguish the fact that while the *value* of the portfolio weight vector will depend on the realization of the state vector, the characterization of the optimal dynamic *policy* will only depend on the time remaining to the horizon, τ .

C.1 Time $t + K - 1$ ($\tau = 1$)

C.1.1 Policy function

The $(t + K - 1)$ objective function remains unchanged relative to the problem we considered in the previous section:

$$\max_{\alpha_{t+K-1}^{(1)}} E_{t+K-1} \left[\exp \{ (1-\gamma) r_{p,t+K} \} \right] \equiv \max_{\alpha_{t+K-1}^{(1)}} \exp \left\{ (1-\gamma) E_{t+K-1} [r_{p,t+K}] + \frac{1}{2} (1-\gamma)^2 \text{Var}_{t+K-1} [r_{p,t+K}] \right\}$$

Using the selector matrix notation the first two moments of the portfolio return are now given by:

$$\begin{aligned} E_{t+K-1} [r_{p,t+K}] &= \alpha_{t+K-1}^{(1)'} \mathbf{H}_x (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-1}) + \mathbf{H}_1 (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-1}) + \frac{1}{2} \alpha_{t+K-1}^{(1)'} (\sigma_x^2 - \Sigma_{xx} \alpha_{t+K-1}^{(1)}) \\ \text{Var}_{t+K-1} [r_{p,t+K}] &= \alpha_{t+K-1}^{(1)'} \Sigma_{xx} \alpha_{t+K-1}^{(1)} + \sigma_1^2 + 2 \alpha_{t+K-1}^{(1)'} \sigma_{1x} \end{aligned}$$

After substituting the moments into the objective function and simplifying we obtain the following first order condition:

$$\mathbf{H}_x (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-1}) + \frac{1}{2} \sigma_x^2 - \Sigma_{xx} \alpha_{t+K-1}^{(1)} + (1-\gamma) (\Sigma_{xx} \alpha_{t+K-1}^{(1)} + \sigma_{1x}) = 0$$

Assuming that the Σ_{xx} is non-singular this equality is solved by:

$$\alpha_{t+K-1}^{(1)} = \frac{1}{\gamma} \Sigma_{xx}^{-1} \left(\mathbf{H}_x (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-1}) + \frac{1}{2} \sigma_x^2 + (1-\gamma) \sigma_{1x} \right) \quad (41)$$

The optimal portfolio can then be written as a affine function of the \mathbf{z}_{t+K-1} vector and can be expressed as:

$$\alpha_{t+K-1}^{(1)} = A_0^{(1)} + A_1^{(1)} \mathbf{z}_{t+K-1} \quad (42)$$

where:

$$A_0^{(1)} = \frac{1}{\gamma} \Sigma_{xx}^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1-\gamma) \sigma_{1x} \right) \quad (43)$$

$$A_1^{(1)} = \frac{1}{\gamma} \Sigma_{xx}^{-1} \Phi_1^x \quad (44)$$

C.1.2 Value function

We begin the derivation of the last period value function by substituting (42) into the expression for the portfolio return moments in order to express them as a function of the $A_i^{(1)}$ coefficients:

$$\begin{aligned} E_{t+K-1} [r_{p,t+K}] &= \left(A_0^{(1)} + A_1^{(1)} \mathbf{z}_{t+K-1} \right)' \left(\Phi_0^x + \Phi_1^x \mathbf{z}_{t+K-1} \right) + \left(\Phi_0^1 + \Phi_1^1 \mathbf{z}_{t+K-1} \right) + \\ &\quad + \frac{1}{2} \left(A_0^{(1)} + A_1^{(1)} \mathbf{z}_{t+K-1} \right)' \left(\sigma_x^2 - \Sigma_{xx} \left(A_0^{(1)} + A_1^{(1)} \mathbf{z}_{t+K-1} \right) \right) \\ &= \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \left(\sigma_x^2 - \Sigma_{xx} A_0^{(1)} \right) \right) + \mathbf{z}_{t+K-1}' A_1^{(1)'} \left(\Phi_0^x + \frac{1}{2} \left(\sigma_x^2 - \Sigma_{xx} A_0^{(1)} \right) \right) + \\ &\quad + \left(\Phi_1^1 + A_0^{(1)'} \left(\Phi_1^x - \frac{1}{2} \Sigma_{xx} A_1^{(1)} \right) \right) \mathbf{z}_{t+K-1} + \mathbf{z}_{t+K-1}' A_1^{(1)'} \left(\Phi_1^x - \frac{1}{2} \Sigma_{xx} A_1^{(1)} \right) \mathbf{z}_{t+K-1} \\ Var_{t+K-1} [r_{p,t+K}] &= \sigma_1^2 + A_0^{(1)'} \left(\Sigma_{xx} A_0^{(1)} + 2\sigma_{1x} \right) + A_0^{(1)'} \Sigma_{xx} A_1^{(1)} \mathbf{z}_{t+K-1} + \\ &\quad + \mathbf{z}_{t+K-1}' A_1^{(1)'} \left(\Sigma_{xx} A_0^{(1)} + 2\sigma_{1x} \right) + \mathbf{z}_{t+K-1}' A_1^{(1)'} \Sigma_{xx} A_1^{(1)} \mathbf{z}_{t+K-1} \end{aligned}$$

Since each summand in the above expressions is a scalar we are free to replace any component with its transpose. This will enable us to gather terms on \mathbf{z}_{t+K-1} and \mathbf{z}_{t+K-1} . Substituting these expressions into the $(t+K-1)$ objective leads to an expression which takes the form of an expectation of an exponential polynomial of the state vector and can be solved in closed form to yield the maximized value function:

$$E_{t+K-1} [\exp \{ (1-\gamma) r_{p,t+K}^* \}] = \exp \left\{ (1-\gamma) \left(B_0^{(1)} + B_1^{(1)} \mathbf{z}_{t+K-1} + \mathbf{z}_{t+K-1}' B_2^{(1)} \mathbf{z}_{t+K-1} \right) \right\} \quad (45)$$

where:

$$B_0^{(1)} \equiv \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) + \frac{1-\gamma}{2} \left(\sigma_1^2 + 2A_0^{(1)'} \sigma_{1x} \right) - \frac{\gamma}{2} A_0^{(1)'} \Sigma_{xx} A_0^{(1)} \quad (46)$$

$$B_1^{(1)} \equiv \Phi_1^1 + A_0^{(1)'} \left(\Phi_1^x - \gamma \Sigma_{xx} A_1^{(1)} \right) + \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1-\gamma) \sigma_{1x} \right)' A_1^{(1)} \quad (47)$$

$$B_2^{(1)} \equiv A_1^{(1)'} \left(\Phi_1^x - \frac{\gamma}{2} \Sigma_{xx} A_1^{(1)} \right) \quad (48)$$

In order to retain the generality of the formulae we refrain from substituting in period specific values for the coefficients of the optimal portfolio (e.g. $A_0^{(1)}$ and $A_1^{(1)}$). By establishing the general form of the relationship between the value and policy function coefficients in adjacent periods we will be able to adapt the solution to a general multi-period setting without having to do any additional derivations.

C.2 Time $t + K - 2$ ($\tau = 2$)

C.2.1 Policy function

The $(t + K - 2)$ problem involves the choice of the portfolio $\alpha_{t+K-2}^{(2)}$ which maximizes the following objective:

$$\frac{W_{t+K-2}^{1-\gamma}}{1-\gamma} E_{t+K-2} \left[\exp \{ (1-\gamma) r_{p,t+K-1} \} \exp \left\{ (1-\gamma) \left(B_0^{(1)} + B_1^{(1)} \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} B_2^{(1)} \mathbf{z}_{t+K-1} \right) \right\} \right] \quad (49)$$

Since the prefactor involving W_{t+K-2} and γ does not affect the optimization problem, we drop it in subsequent expressions in order to economize on notation. We begin the derivation of the optimal policy function by making use of the Campbell-Viceira approximation to express the time $(t + K - 1)$ log portfolio return in terms of the \mathbf{z}_{t+K-1} . This allows us to collect terms in the exponential and compute its expectation.

The $(t + K - 1)$ log portfolio return can be approximated by:

$$\begin{aligned} r_{p,t+K-1} &= \alpha_{t+K-2}^{(2)'} \mathbf{x}_{t+K-1} + r_{1,t+K-1} + \frac{1}{2} \alpha_{t+K-2}^{(2)'} \left(\sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} \right) \\ &= \begin{bmatrix} 1 & \alpha_{t+K-2}^{(2)} & \mathbf{0} \end{bmatrix} \mathbf{z}_{t+K-1} + \frac{1}{2} \alpha_{t+K-2}^{(2)'} \left(\sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} \right) \end{aligned}$$

which - after substituting into the objective function and dropping the prefactor - yields:

$$\begin{aligned} E_{t+K-2} \left[\exp \left\{ (1-\gamma) \left(\frac{1}{2} \alpha_{t+K-2}^{(2)'} \left(\sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} \right) + B_0^{(1)} + \left(B_1^{(1)} + \begin{bmatrix} 1 & \alpha_{t+K-2}^{(2)} & \mathbf{0} \end{bmatrix} \right) \mathbf{z}_{t+K-1} + \right. \right. \\ \left. \left. + \mathbf{z}'_{t+K-1} B_2^{(1)} \mathbf{z}_{t+K-1} \right) \right\} \right] \quad (50) \end{aligned}$$

From the point of view of the investor at time $(t + K - 2)$ only the last two summands in the exponential are random. The computation of the expectation can therefore be simplified by focusing on the random components. We therefore focus on evaluating the following expression given the investor's choice of $\alpha_{t+K-2}^{(2)}$:

$$E_{t+K-2} \left[\exp \left\{ (1-\gamma) \left(\left(B_1^{(1)} + \begin{bmatrix} 1 & \alpha_{t+K-2}^{(2)} & \mathbf{0} \end{bmatrix} \right) \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} B_2^{(1)} \mathbf{z}_{t+K-1} \right) \right\} \right] \quad (51)$$

The result of this computation is then substituted back into (50) to obtain the expectation of the value function conditional on the investor's portfolio choice decision, which is then maximized over $\alpha_{t+K-2}^{(2)}$.

Since the distribution of \mathbf{z}_{t+K-1} conditional on time $(t + K - 2)$ information is :

$$\mathbf{z}_{t+K-1|t+K-2} \sim N(\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}, \Sigma_v)$$

it is clear that the evaluation of (51) is tantamount to computing the expectation of an exponential of a quadratic polynomial in a normal random variable. We can proceed from here in one of two ways depending on whether we use the continuous time approximation to (51) or not. We present the solutions for both scenarios, but we elect to focus on the continuous-time approximation since this approach is complementary to the continuous-time approximation to the portfolio return used throughout the paper.

In order to economize on space we first introduce the following notation:

$$\begin{aligned} C_1^{(2)} &\equiv (1-\gamma) \begin{bmatrix} 1 \\ \alpha_{t+K-2}^{(2)} \\ \mathbf{0} \end{bmatrix} + (1-\gamma) B_1^{(1)'} \\ C_2^{(2)} &\equiv (1-\gamma) B_2^{(1)} \end{aligned}$$

which, when combined with the conditional distribution of $\mathbf{z}_{t+K-1|t+K-2}$, allows us to re-express (51) as:

$$E_{t+K-2} \left[\exp \left\{ C_1^{(2)'} \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} C_2^{(2)} \mathbf{z}_{t+K-1} \right\} \right] = E_{t+K-2} \left[\exp \left\{ C_1^{(2)'} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2} + \mathbf{v}_{t+K-2}) + \right. \right. \quad (52)$$

$$\left. \left. + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2} + \mathbf{v}_{t+K-2})' C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2} + \mathbf{v}_{t+K-2}) \right\} \right]$$

Inspection of this equation makes immediately clear that the terms appearing in the exponential are either constants, linear combinations of the \mathbf{v}_{t+K-2} shocks or quadratic forms in \mathbf{v}_{t+K-2} (i.e. the term $\mathbf{v}'_{t+K-2} C_2^{(2)} \mathbf{v}_{t+K-2}$). We can further reorder the terms in equation (52) as follows

$$\begin{aligned} E_{t+K-2} [\exp(\cdot)] &= \exp \left(C_1^{(2)'} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) \right) \\ &\quad \times E_{t+K-2} \left[\exp \left\{ C_1^{(2)'} \mathbf{v}_{t+K-2} + \mathbf{v}'_{t+K-2} C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' C_2^{(2)} \mathbf{v}'_{t+K-2} + \right. \right. \\ &\quad \left. \left. + \mathbf{v}'_{t+K-2} C_2^{(2)} \mathbf{v}_{t+K-2} \right\} \right] \\ &= \exp \left(C_1^{(2)'} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) \right) \\ &\quad \times E_{t+K-2} \left[\exp \left\{ C_1^{(2)'} \mathbf{v}_{t+K-2} + \mathbf{v}'_{t+K-2} C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' C_2^{(2)} \mathbf{v}_{t+K-2} \right\} \right. \\ &\quad \left. \times \exp \left\{ \mathbf{v}'_{t+K-2} C_2^{(2)} \mathbf{v}_{t+K-2} \right\} \right] \\ &= \exp \left(C_1^{(2)'} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) \right) \\ &\quad \times E_{t+K-2} \left[\exp \left\{ \mathbf{v}'_{t+K-2} C_2^{(2)} \mathbf{v}_{t+K-2} \right\} \times \exp \left\{ \left(C_1^{(2)'} + 2(\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' \tilde{C}_2^{(2)} \right) \mathbf{v}_{t+K-2} \right\} \right], \quad (53) \end{aligned}$$

where we have introduced the convention that *any matrix annotated with a tilde is to be interpreted as the average of the non-tilde matrix and its transpose*. For example:

$$\tilde{C}_2^{(2)} \equiv \frac{1}{2} \left(C_2^{(2)} + C_2^{(2)'} \right)$$

Clearly, all such matrices are symmetric by construction. The tilde convention is introduced to ensure that the validity of our results does not hinge on the symmetry of the underlying matrices - in this case $C_2^{(2)}$. However, whenever we show that the underlying matrix is symmetric we drop the tilde notation to avoid redundancy since in those cases the underlying-matrix and the tilde-matrix are identical.

There are two approaches that we can adopt to evaluate the expectation operator in equation (53). First, we can introduce an additional approximation which is exact in continuous time and - in this sense - complements the portfolio return approximation; second, we can simply evaluate the expectation by "completing the square" of the expression in the exponential, following CCV (2003). We present both results for completeness even though they provide empirically indistinguishable results with return data measured at the monthly frequency.

Complementary approximation

We begin by presenting the complementary approximation, which we utilize to obtain the value of the first expectation. To arrive at our result we simply rely on the fact that all terms of order exceeding dt have vanishing probability limits as the time interval shrinks to zero. Thus in the limit of continuous time, *only* the linear terms in equation (53) remain random. More precisely, the linear terms are of order dZ , while the quadratic terms are of order dt (i.e. asymptotically they are deterministic). Thus

$$\lim_{\Delta t \rightarrow 0} \exp \left\{ \mathbf{v}'_{t+K-2} C_2^{(2)} \mathbf{v}_{t+K-2} \right\} / dt = \exp \left\{ \text{tr}(C_2^{(2)} \Sigma_v) \right\},$$

and equation (53) becomes

$$\begin{aligned} & \exp \left(C_1^{(2)'} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + \text{tr}(C_2^{(2)} \Sigma_v) \right) \\ & \times E_{t+K-2} \left[\exp \left\{ \left(C_1^{(2)'} + 2 (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' \tilde{C}_2^{(2)} \right) \mathbf{v}_{t+K-2} \right\} \right] \end{aligned} \quad (54)$$

Fortunately, the expectation in (54) is nothing more than an expectation of a log-normal random variable, and is given by

$$\begin{aligned} & E_{t+K-2} \left[\exp \left\{ \left(C_1^{(2)'} + 2 (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' \tilde{C}_2^{(2)} \right) \mathbf{v}_{t+K-2} \right\} \right] = \\ & = \exp \left\{ \frac{1}{2} \left(C_1^{(2)'} + 2 (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' \tilde{C}_2^{(2)} \right)' \Sigma_v \left(C_1^{(2)'} + 2 (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' \tilde{C}_2^{(2)} \right) \right\} \end{aligned} \quad (55)$$

Combining (54) and (55) we see that in the continuous-time limit (51) is equal to:

$$\lim_{\Delta t \rightarrow 0} E_{t+K-2} \left[\exp \left\{ C_1^{(2)'} \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} C_2^{(2)} \mathbf{z}_{t+K-1} \right\} \right] = \exp \left\{ \text{tr}(C_2^{(2)} \Sigma_v) + D_0^{(t+K-2)} + \frac{1}{2} D_1^{(t+K-2)} \Sigma_v^{-1} D_1^{(t+K-2)'} \right\} \quad (56)$$

where

$$\begin{aligned} D_0^{(2)} &= C_1^{(2)'} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) + (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' C_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) \\ D_1^{(2)} &= C_1^{(2)'} + 2 (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2})' \tilde{C}_2^{(2)}. \end{aligned}$$

Exact evaluation

Had we elected not to rely on the continuous-time approximation to the exponential quadratic term a closed-form solution would still be available. In particular, the results derived in CCV (2003) provide a direct route to the computing the expectation in (51). Using the previously introduced notation we can apply equation (46) from the technical appendix to CCV (2003) to arrive at the following expression:

$$\begin{aligned} & E_{t+K-2} \left[\exp \left\{ C_1^{(2)'} \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} C_2^{(2)} \mathbf{z}_{t+K-1} \right\} \right] \\ & = \frac{|\Sigma_v|^{-\frac{1}{2}}}{\left| \left(\Sigma_v^{-1} - 2D_2^{(t+K-2)} \right) \right|^{\frac{1}{2}}} \exp \left(D_0^{(t+K-2)} + \frac{1}{2} D_1^{(t+K-2)} \left(\Omega^{(1)} \right)^{-1} D_1^{(t+K-2)'} \right) \end{aligned} \quad (57)$$

where we have additionally introduced $D_2^{(2)} = C_2^{(2)}$ and

$$\Omega^{(1)} \equiv \left(\Sigma_v^{-1} - 2D_2^{(2)} \right)^{-1} = \left(\Sigma_v^{-1} - 2(1-\gamma)B_2^{(1)} \right)^{-1}. \quad (58)$$

The similarities and differences between equation (56) and equation (57) are immediate.

Optimal portfolio policy

With (56) or (57) in hand, we can replace the (conditional) expectation of the random component of the value function, (51), with a function that only depends on $\alpha_{t+K-2}^{(2)}$ and the time $(t+K-2)$ state vector. Since equations (56) and (57) are functionally identical, the derivation and the form of the solution under each is identical—and numerically indistinguishable in our empirical calibration exercise with return data measured at the monthly frequency. We now derive the solution implied by (57), and state the solution implied by (56).

To derive the optimal portfolio rule, we economize on notation by dropping the prefactor arising from the computation of the expectation of the exponential polynomial (57), since constant prefactors do not affect the solution to the maximization problem.²⁸ This allows us to reformulate the problem facing the investor at time $(t + K - 2)$ as:

$$\max_{\alpha_{t+K-2}^{(2)}} \left\{ (1 - \gamma) \left(\frac{1}{2} \alpha_{t+K-2}^{(2)'} \left(\sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} \right) + B_0^{(1)} + D_0^{(t+K-2)} + \frac{1}{2} D_1^{(t+K-2)} \Omega^{(1)} D_1^{(t+K-2)'} \right) \right\}$$

Taking first derivatives with respect to the vector of portfolio weights we arrive at the following first order condition for $\alpha_{t+K-2}^{(2)}$:

$$\begin{aligned} 0 &= (1 - \gamma) \left(\Phi_0^x + \Phi_1^x \mathbf{z}_{t+K-2} + \frac{1}{2} \sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} \right) + \frac{1}{2} \left(\frac{\partial D_1^{(t+K-2)'}}{\partial \alpha_{t+K-2}} \right)' \left(\Omega^{(1)} + \Omega^{(1)'} \right) D_1^{(t+K-2)'} \\ &= (1 - \gamma) \left(\Phi_0^x + \Phi_1^x \mathbf{z}_{t+K-2} + \frac{1}{2} \sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} \right) + \left(\frac{\partial D_1^{(t+K-2)'}}{\partial \alpha_{t+K-2}} \right)' \tilde{\Omega}^{(1)} D_1^{(t+K-2)'} \end{aligned}$$

Given the definition for $D_1^{(2)}$, we can relate its derivative with respect to the portfolio allocation to the risky asset excess return selector matrix (i.e. the selector matrix which returns the rows denoted by \mathbf{x}):

$$\frac{\partial D_1^{(t+K-2)'}}{\partial \alpha_{t+K-2}} = (1 - \gamma) \begin{bmatrix} \mathbf{0} \\ \mathbf{I}_{n-1} \\ \mathbf{0} \end{bmatrix} \equiv \left(\frac{\partial D_1^{(t+K-2)'}}{\partial \alpha_{t+K-2}} \right)' = (1 - \gamma) \mathbf{H}_x$$

This ultimately leads us to the following form of the first order condition:

$$\begin{aligned} 0 &= \Phi_0^x + \Phi_1^x \mathbf{z}_{t+K-2} + \frac{1}{2} \sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} + \mathbf{H}_x \tilde{\Omega}^{(1)} \left(C_1^{(2)} + 2\tilde{C}_2^{(2)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) \right) \\ &= \Phi_0^x + \Phi_1^x \mathbf{z}_{t+K-2} + \frac{1}{2} \sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} + (1 - \gamma) \tilde{\Omega}_x^{(1)} \left(\begin{bmatrix} 1 \\ \alpha_{t+K-2} \\ \mathbf{0} \end{bmatrix} + B_1^{(1)'} + 2\tilde{B}_2^{(1)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) \right) \\ &= \Phi_0^x + \Phi_1^x \mathbf{z}_{t+K-2} + \frac{1}{2} \sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} + (1 - \gamma) \left(\tilde{\Omega}_{x1}^{(1)} + \tilde{\Omega}_{xx}^{(1)} \alpha_{t+K-2}^{(2)} + \tilde{\Omega}_x B_1^{(1)'} + \right. \\ &\quad \left. + 2\tilde{\Omega}_x^{(1)} \tilde{B}_2^{(1)} (\Phi_0 + \Phi_1 \mathbf{z}_{t+K-2}) \right) \\ &= \left((1 - \gamma) \tilde{\Omega}_{xx}^{(1)} - \Sigma_{xx} \right) \alpha_{t+K-2}^{(2)} + \Phi_0^x + (1 - \gamma) \left(\tilde{\Omega}_{x1}^{(1)} + \tilde{\Omega}_x^{(1)} \left(B_1^{(1)'} + 2\tilde{B}_2^{(1)} \Phi_0 \right) \right) + \\ &\quad + \frac{1}{2} \sigma_x^2 + \left(\Phi_1^x + 2(1 - \gamma) \tilde{\Omega}_x^{(1)} \tilde{B}_2^{(1)} \Phi_1 \right) \mathbf{z}_{t+K-2} \end{aligned}$$

Solving the first order condition yields an expression for the optimal portfolio holding at time $(t + K - 2)$ in the form of an affine function of \mathbf{z}_{t+K-2} :

$$\alpha_{t+K-2}^{(2)} = A_0^{(2)} + A_1^{(2)} \mathbf{z}_{t+K-2} \quad (59)$$

where:

$$A_0^{(2)} = \left(\Sigma_{xx} - (1 - \gamma) \tilde{\Omega}_{xx}^{(1)} \right)^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1 - \gamma) \left(\tilde{\Omega}_{x1}^{(1)} + \tilde{\Omega}_x^{(1)} \left(B_1^{(1)'} + 2\tilde{B}_2^{(1)} \Phi_0 \right) \right) \right) \quad (60)$$

$$A_1^{(2)} = \left(\Sigma_{xx} - (1 - \gamma) \tilde{\Omega}_{xx}^{(1)} \right)^{-1} \left(\Phi_1^x + 2(1 - \gamma) \tilde{\Omega}_x^{(1)} \tilde{B}_2^{(1)} \Phi_1 \right) \quad (61)$$

²⁸While dropping the prefactors arising from the evaluation of the expectation does not affect the choice of the optimal portfolio, it will affect the value function for the problem. By keeping track of these prefactors it is possible to derive a closed-form expression for the value function, and consequently, conduct an extensive welfare analysis.

Similar computations based on (56) lead to the following expression:

$$A_0^{(2)} = \frac{1}{\gamma} \Sigma_{xx}^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1-\gamma) \left(\Sigma_{x1} + \Sigma_x \left(B_1^{(1)'} + 2\tilde{B}_2^{(1)} \Phi_0 \right) \right) \right) \quad (62)$$

$$A_1^{(2)} = \frac{1}{\gamma} \Sigma_{xx}^{-1} \left(\Phi_1^x + 2(1-\gamma) \Sigma_x \tilde{B}_2^{(1)} \Phi_1 \right) \quad (63)$$

It is important to note that (62) and (63) are identical to (60) and (61), except that $\tilde{\Omega}^{(1)}$ is replaced with Σ_v .

C.2.2 Value function

In order to complete the derivation of the general recursive solution we still need to develop the relevant expressions for the coefficients in the maximized value function, $B_i^{(2)}$. This will enable us to determine the relationship with the preceding values of the $A_i^{(1)}$ and $B_i^{(1)}$ coefficients. Substituting the optimal portfolio allocation for $(t+K-2)$ into the value function, (49), and omitting any prefactors multiplying the expectation we obtain:

$$\begin{aligned} & E_{t+K-2} \left[\exp \{ (1-\gamma) r_{p,t+K-1} \} \exp \left\{ (1-\gamma) \left(B_0^{(1)} + B_1^{(1)} \mathbf{z}_{t+K-1} + \mathbf{z}'_{t+K-1} B_2^{(1)} \mathbf{z}_{t+K-1} \right) \right\} \right] = \\ & = \exp \left\{ (1-\gamma) \left(\frac{1}{2} \alpha_{t+K-2}^{(2)'} \left(\sigma_x^2 - \Sigma_{xx} \alpha_{t+K-2}^{(2)} \right) + B_0^{(1)} \right) + D_0^{(t+K-2)} + \frac{1}{2} D_1^{(t+K-2)} \Omega^{(1)} D_1^{(t+K-2)'} \right\} \\ & = \exp \left\{ (1-\gamma) \left(B_0^{(2)} + B_1^{(2)} \mathbf{z}_{t+K-2} + \mathbf{z}'_{t+K-2} B_2^{(2)} \mathbf{z}_{t+K-2} \right) \right\} \end{aligned} \quad (64)$$

where:

$$\begin{aligned} B_0^{(2)} & = \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) + \frac{1-\gamma}{2} \left(\Omega_{11}^{(1)} + 2A_0^{(1)'} \tilde{\Omega}_{x1}^{(1)} \right) - \\ & \quad - \frac{1}{2} A_0^{(1)'} \left(\Sigma_{xx} - (1-\gamma) \Omega_{xx}^{(1)} \right) A_0^{(1)} + B_0^{(1)} + \left(B_1^{(1)} + \Phi_0' B_2^{(1)} \right) \Phi_0 + \\ & \quad + (1-\gamma) B_1^{(1)} \left(\tilde{\Omega}_1^{(1)'} + \frac{1}{2} \Omega^{(1)} B_1^{(1)'} + \tilde{\Omega}_x^{(1)'} A_0^{(1)} \right) + \\ & \quad + (1-\gamma) \Phi_0' \left(2\Lambda^{(1)} \Phi_0 + \Xi_1^{(1)'} + \Xi_x^{(1)'} A_0^{(1)} + \Gamma^{(1)} B_1^{(1)'} \right) \end{aligned} \quad (65)$$

$$\begin{aligned} B_1^{(2)} & = \Phi_1^1 + A_0^{(1)'} \left(\Phi_1^x - \left(\Sigma_{xx} - (1-\gamma) \tilde{\Omega}_{xx}^{(1)} \right) A_1^{(1)} \right) + \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1-\gamma) \tilde{\Omega}_{x1}^{(1)} \right)' A_1^{(1)} + \\ & \quad + \left(B_1^{(1)} + 2\Phi_0' \tilde{B}_2^{(1)} \right) \Phi_1 + (1-\gamma) \left(B_1^{(1)} \tilde{\Omega}_x^{(1)'} + \Phi_0' \Xi_x^{(1)'} \right) A_1^{(1)} + \\ & \quad + (1-\gamma) \left(4\Phi_0' \tilde{\Lambda}^{(1)} + \Xi_1^{(1)} + A_0^{(1)'} \Xi_x^{(1)} + B_1^{(1)} \Xi^{(1)} \right) \Phi_1 \end{aligned} \quad (66)$$

$$\begin{aligned} B_2^{(2)} & = A_1^{(1)'} \left(\Phi_1^x - \frac{1}{2} \left(\Sigma_{xx} - (1-\gamma) \Omega_{xx}^{(1)} \right) A_1^{(1)} \right) + \Phi_1' \left(B_2^{(1)} + 2(1-\gamma) \Lambda^{(1)} \right) \Phi_1 + \\ & \quad + (1-\gamma) \Phi_1' \Xi_x^{(1)'} A_1^{(1)} \end{aligned} \quad (67)$$

where we have defined the auxiliary matrices $\Lambda^{(1)} = \tilde{B}_2^{(1)} \Omega^{(1)} \tilde{B}_2^{(1)'}$, $\Gamma^{(1)} = 2\tilde{B}_2^{(1)} \tilde{\Omega}^{(1)}$ and $\Xi^{(1)} = \Gamma^{(1)'}$.

In the limit of continuous-time we obtain a set of simplifications analogous to those appearing in the derivation of the policy function. In particular, making the relevant substitutions we find the following

expressions for the value function coefficients:

$$B_0^{(2)} = \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) + \frac{1-\gamma}{2} \left(\Sigma_{11} + 2A_0^{(1)'} \Sigma_{x1} \right) - \frac{\gamma}{2} A_0^{(1)'} \Sigma_{xx} A_0^{(1)} + B_0^{(1)} + \left(B_1^{(1)} + \Phi_0' B_2^{(1)} \right) \Phi_0 + (1-\gamma) B_1^{(1)} \left(\Sigma_1' + \frac{1}{2} \Sigma B_1^{(1)'} + \Sigma_x' A_0^{(1)} \right) + (1-\gamma) \Phi_0' \left(2\Lambda^{(1)} \Phi_0 + \Xi_1^{(1)'} + \Xi_x^{(1)'} A_0^{(1)} + \Gamma^{(1)} B_1^{(1)'} \right) \quad (68)$$

$$B_1^{(2)} = \Phi_1^1 + A_0^{(1)'} \left(\Phi_1^x - \gamma \Sigma_{xx} \right) A_1^{(1)} + \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1-\gamma) \Sigma_{x1} \right)' A_1^{(1)} + \left(B_1^{(1)} + 2\Phi_0' \tilde{B}_2^{(1)} \right) \Phi_1 + (1-\gamma) \left(B_1^{(1)} \Sigma_x' + \Phi_0' \Xi_x^{(1)'} \right) A_1^{(1)} + (1-\gamma) \left(4\Phi_0' \tilde{\Lambda}^{(1)} + \Xi_1^{(1)} + A_0^{(1)'} \Xi_x^{(1)} + B_1^{(1)} \Xi^{(1)} \right) \Phi_1 \quad (69)$$

$$B_2^{(2)} = A_1^{(1)'} \left(\Phi_1^x - \frac{\gamma}{2} \Sigma_{xx} A_1^{(1)} \right) + \Phi_1' \left(B_2^{(1)} + 2(1-\gamma) \Lambda^{(1)} \right) \Phi_1 + (1-\gamma) \Phi_1' \Xi_x^{(1)'} A_1^{(1)} \quad (70)$$

where we have defined the auxiliary matrices $\Lambda^{(1)} = \tilde{B}_2^{(1)} \Sigma_v \tilde{B}_2^{(1)'}$, $\Gamma^{(1)} = 2\tilde{B}_2^{(1)} \Sigma_v$ and $\Xi^{(1)} = \Gamma^{(1)'}$.

C.2.3 The general recursive solution

Since we have refrained from substituting any period specific values in the derivation of the time $(t+K-2)$ formulas of the policy and value function coefficients, the results of the previous section provide the complete characterization of the link between the policy and value functions in any two adjacent periods. In other words, the link between the coefficients governing policy and value functions in periods (τ) and $(\tau-1)$, is exactly the same as the link between periods (2) and (1). Therefore the general recursive solution is obtained simply by making the following replacements:

$$(1) \rightarrow (\tau-1) \quad \text{and} \quad (2) \rightarrow (\tau)$$

in the recursive formulas for the policy function (eqs. (60) and (61)) and value function coefficients (eqs. (65), (66) and (67)). The base cases for the recursions in the policy and value function coefficients are given by, (eqs. (43) and (44)) and (eqs. (46), (47) and (48)), respectively. These two sets of recursions, along with their base cases, provide a complete characterization of the optimal dynamically consistent portfolio policy.

C.2.4 Matrix symmetry

Having derived the general solution we now check whether any further simplification is possible as a result of the symmetry of the matrices in the policy and value function recursions. In particular, throughout the derivations we used the *tilde convention* to assure that our results would be robust to any underlying matrix non-symmetries. Here we verify that the tilde notation can indeed be dropped.

The matrices which were affected by the tilde notation are $B_2^{(\tau)}$, $C_2^{(\tau)}$ and $\Omega^{(1)}$. However, as can be seen directly from the derivation the second two matrices are merely transformations of the corresponding $B_2^{(\tau)}$ matrices. Therefore if we can ascertain that $B_2^{(\tau)}$ is symmetric then so will $C_2^{(\tau)}$ and $\Omega^{(1)}$. In order to ascertain this fact we need to show that $B_2^{(1)}$ is symmetric and the recursion relating $B_2^{(\tau)}$ to $B_2^{(\tau-1)}$ preserves this symmetry. The first fact is readily established since:

$$B_2^{(1)} \equiv A_1^{(1)'} \left(\Phi_1^x - \frac{\gamma}{2} \Sigma_{xx} A_1^{(1)} \right) = \frac{1}{2\gamma} \Phi_1^{x'} \Sigma_{xx} \Phi_1^x$$

which is clearly symmetric.

C.3 Special cases

To gain additional intuition regarding the solution and to confirm its validity we examine two corner cases: $\gamma = 1$ and $\gamma \rightarrow \infty$. Below we present the simplified forms of the base cases and recursive relationships:

1. **Myopic portfolio choice.** Substituting in $\gamma = 1$ into the general relationships derived in the previous two sections results in the following expressions. The portfolio choice policy coefficients evolve according to:

$$\begin{aligned} A_0^{(\tau)} &= \Sigma_{xx}^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) \quad \forall \tau \\ A_1^{(\tau)} &= \Sigma_{xx}^{-1} \Phi_1^x \quad \forall \tau \end{aligned}$$

since:

$$\Omega^{(1)} = \Sigma_v \quad \forall \tau$$

Consequently, the base case for the policy function recursion becomes:

$$\begin{aligned} B_0^{(1)} &= \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 - \frac{1}{2} \Sigma_{xx} A_0^{(1)} \right) = \Phi_0^1 \\ B_1^{(1)} &= \Phi_1^1 + A_0^{(1)'} \left(\Phi_1^x - \Sigma_{xx} A_1^{(1)} \right) + \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right)' A_1^{(1)} = \Phi_1^1 + \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right)' \Sigma_{xx}^{-1} \Phi_1^x \\ B_2^{(1)} &= A_1^{(1)'} \left(\Phi_1^x - \frac{1}{2} \Sigma_{xx} A_1^{(1)} \right) = \frac{1}{2} \Phi_1^{x'} \Sigma_{xx}^{-1} \Phi_1^x \end{aligned}$$

and the recursive relationship itself is given by:

$$\begin{aligned} B_0^{(\tau)} &= \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 - \frac{1}{2} \Sigma_{xx} A_0^{(\tau-1)} \right) + B_0^{(\tau-1)} + \left(B_1^{(\tau-1)} + \Phi_0' B_2^{(\tau-1)} \right) \Phi_0 \\ &= \Phi_0^1 + \frac{1}{2} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right)' \Sigma_{xx}^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) + B_0^{(\tau-1)} + \left(B_1^{(\tau-1)} + \Phi_0' B_2^{(\tau-1)} \right) \Phi_0 \\ B_1^{(\tau)} &= \Phi_1^1 + A_0^{(\tau-1)'} \left(\Phi_1^x - \Sigma_{xx} A_1^{(\tau-1)} \right) + \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right)' A_1^{(\tau-1)} + \left(B_1^{(\tau-1)} + 2\Phi_0' \tilde{B}_2^{(\tau-1)} \right) \Phi_1 \\ &= \Phi_1^1 + \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right)' \Sigma_{xx}^{-1} \Phi_1^x + \left(B_1^{(\tau-1)} + 2\Phi_0' \tilde{B}_2^{(\tau-1)} \right) \Phi_1 \\ B_2^{(\tau)} &= A_1^{(\tau-1)'} \left(\Phi_1^x - \frac{1}{2} \Sigma_{xx} A_1^{(\tau-1)} \right) + \Phi_1' B_2^{(\tau-1)} \Phi_1 = \frac{1}{2} \Phi_1^{x'} \Sigma_{xx}^{-1} \Phi_1^x + \Phi_1' B_2^{(\tau-1)} \Phi_1 \end{aligned}$$

2. **Infinite risk aversion** ($\gamma \rightarrow \infty$). In order to examine the asymptotic behavior of the recursive relationships determining portfolio choice it is useful to first focus on the order of the terms in γ , $O(\gamma^{-n})$. Terms for which $n \geq 1$ will disappear in the limit, those with $n = 0$ will yield constants, and those with $n < 0$ will diverge. Intuitively, we expect that while terms in the value function recursion may diverge those in the portfolio policy functions will not. Since the value function is an exponential polynomial function, divergence to negative infinity is permissible as it indicates that the value of the problem tends to zero.

We begin with the base case for the policy function recursion. It is immediate that the terms, $A_0^{(1)}$, and $A_1^{(1)}$, are of order $O(1)$ and $O(\gamma^{-1})$. In particular we obtain the following limits:

$$A_0^{(1)} \rightarrow -\Sigma_{xx}^{-1}\sigma_{1x} \quad A_1^{(1)} \rightarrow 0$$

A similar examination for the base case of the policy function recursion reveals that, $B_0^{(1)}$ has order $O(\gamma)$ and diverges to negative infinity, $B_1^{(1)}$ has order $O(1)$ and converges to:

$$B_1^{(1)} \rightarrow \Phi_1^1 - \sigma'_{1x}\Sigma_{xx}^{-1}\Phi_1^x$$

Lastly, $B_2^{(1)}$ is of order $O(\gamma^{-1})$ and hence converges to zero. With these results we can infer that $\Omega^{(1)}$ will be $O(1)$ in the coefficient of risk aversion and, consequently, each iteration will lead to a non-negligible adjustment to the covariance matrix of the shocks. In particular we have that:

$$\Omega^{(1)} \rightarrow \left(\Sigma_v^{-1} + \Phi_1^{x'}\Sigma_{xx}^{-1}\Phi_1^x\right)^{-1}$$

Next we proceed to the $t+K-2$ policy function coefficients. To make the analysis easier to interpret we can write the expressions in the following form:

$$\begin{aligned} A_0^{(2)} &= \left(\frac{1}{\gamma}\Sigma_{xx} - \left(\frac{1}{\gamma} - 1\right)\tilde{\Omega}_{xx}^{(1)}\right)^{-1} \left(\frac{1}{\gamma}\Phi_0^x + \frac{1}{2\gamma}\sigma_x^2 + \left(\frac{1}{\gamma} - 1\right)\left(\tilde{\Omega}_{x1}^{(1)} + \right. \right. \\ &\quad \left. \left. + \tilde{\Omega}_x^{(1)}\left(B_1^{(1)'} + 2\tilde{B}_2^{(1)}\Phi_0\right)\right)\right) \\ A_1^{(2)} &= \left(\frac{1}{\gamma}\Sigma_{xx} - \left(\frac{1}{\gamma} - 1\right)\tilde{\Omega}_{xx}^{(1)}\right)^{-1} \left(\frac{1}{\gamma}\Phi_1^x + 2\left(\frac{1}{\gamma} - 1\right)\tilde{\Omega}_x^{(1)}\tilde{B}_2^{(1)}\Phi_1\right) \end{aligned}$$

Taking limits and noting the previously derived orders of the constituent terms leads to:

$$\begin{aligned} A_0^{(2)} &\rightarrow -\left(\tilde{\Omega}_{xx}^{(1)}\right)^{-1}\left(\tilde{\Omega}_{x1}^{(1)} + \tilde{\Omega}_x^{(1)}B_1^{(1)'}\right) \\ A_1^{(2)} &\rightarrow 0 \end{aligned}$$

Continuing in this fashion it is possible to show that the order of each of the terms appearing in the solution is conserved under the recursion, and therefore, that the results we just derived apply to all time periods. In particular - with a slight abuse of limit notation - the recursion for the value function coefficients truncates to:

$$\begin{aligned} B_0^{(2)} &\rightarrow -\infty \\ B_1^{(2)} &\rightarrow \Phi_1^1 + A_0^{(1)'}\Phi_1^x + B_1^{(1)}\Phi_1 - \gamma B_1^{(1)}\tilde{\Omega}_x^{(1)'}A_1^{(1)} - \\ &\quad - \gamma\left(\Xi_1^{(1)} + A_0^{(1)'}\Xi_x^{(1)} + B_1^{(1)}\Xi^{(1)}\right)\Phi_1 \\ B_2^{(2)} &\rightarrow 0 \end{aligned}$$

In particular, although some expressions continue to have γ in them they are $O(1)$ due to the presence of offsetting terms that are $O(\gamma^{-1})$, which are not fully expanded. For completeness we note that in arriving at these simplified expressions we have made use of the fact that $\Lambda^{(1)} \sim O(\gamma^{-2})$, $\Gamma^{(1)} \sim O(\gamma^{-1})$ and $\Xi^{(1)} \sim O(\gamma^{-1})$, and hence all converge to zero as γ diverges.

Consequently, it is clear that as the agent becomes infinitely risk averse the policy function becomes decoupled from the intercept vector, Φ_0 , of the vector autoregressive model. This result is similar to the optimal portfolio allocation of an infinitely risk averse investor maximizing a mean-variance objective based on K-period moments. In fact, the portfolio allocation at $t+K-1$ is precisely the global minimum variance portfolio for the one-period problem.

3. **No Predictability** ($\Phi_1 = \mathbf{0}$) When returns are unpredictable the slope matrix of the VAR(1) model is identically equal to zero. In this case the first period policy function coefficients truncate to:

$$\begin{aligned} A_0^{(1)} &= \frac{1}{\gamma} \Sigma_{xx}^{-1} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 + (1 - \gamma) \sigma_{1x} \right) \\ A_1^{(1)} &= \mathbf{0} \end{aligned}$$

and lead to the following set of coefficients for the base case of the value function recursion:

$$\begin{aligned} B_0^{(1)} &= \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) + \frac{1 - \gamma}{2} \left(\sigma_1^2 + 2A_0^{(1)'} \sigma_{1x} \right) - \frac{\gamma}{2} A_0^{(1)'} \Sigma_{xx} A_0^{(1)} \\ B_1^{(1)} &= B_2^{(1)} = 0 \end{aligned}$$

Consequently, it is easy to see that we have $\Omega^{(1)} = \Sigma_v$ which, along with the values of the value function coefficients, yields:

$$A_0^{(2)} = A_0^{(1)} \quad A_1^{(2)} = A_0^{(1)}$$

Continuing, by substituting these values into the time (2) value function coefficients leads to:

$$\begin{aligned} B_0^{(2)} &= \Phi_0^1 + A_0^{(1)'} \left(\Phi_0^x + \frac{1}{2} \sigma_x^2 \right) + \frac{1 - \gamma}{2} \left(\sigma_1^2 + 2A_0^{(1)'} \sigma_{x1} \right) - \frac{\gamma}{2} A_0^{(1)'} \Sigma_{xx} A_0^{(1)} + B_0^{(1)} \\ B_1^{(2)} &= B_2^{(2)} = 0 \end{aligned}$$

since $\Lambda^{(1)} = \Gamma^{(1)} = \Xi^{(1)} = \mathbf{0}$. Given that the portfolio rule is only affected by the B_1 and B_2 matrices and the fact that the value of these two matrices only depends on their own lagged values and the lagged value of A_1 , one can immediately conclude that the portfolio allocation is constant across time and uniquely determined by A_0 . The value of the problem, on the other hand, is determined by B_0 .

D Appendix: Temporal Aggregation of the VAR(1) model

The VAR(1) class of time-series models is closed under temporal aggregation. In other words, if \mathbf{z}_{t+1} follows an VAR(1) process, so will $\sum_{i=0}^n \mathbf{z}_{t+1+i}$. This observation enables us to interpret our model as applying at different data frequencies, and consequently, allows us to examine the effect of the re-balancing frequency on the agent's welfare.

Consider the constrained version of the VAR(1) model in which the state variables forecast themselves and returns, but the returns are forecast only by the state variables and do not have any forecasting power themselves. This system is equivalent to the system of equations:

$$\begin{aligned}\mathbf{r}_{t+1} &= \Phi_{0r} + \Phi_{1r}\mathbf{s}_t + \mathbf{v}_{r,t+1} \\ \mathbf{s}_{t+1} &= \Phi_{0s} + \Phi_{1s}\mathbf{s}_t + \mathbf{v}_{s,t+1}\end{aligned}$$

The covariance matrix of the shocks to the constrained system can be decomposed into three matrices, Σ_{rr} , Σ_{rs} and Σ_{ss} .

Now consider forecasting k-period returns at time t using the time t values of the state variables. A simple recursion indicates that:

$$\sum_{i=0}^{k-1} \mathbf{r}_{t+1+i} = k \cdot \Phi_{0,r} + \Phi_{1,r} \cdot \sum_{j=0}^{k-1} \mathbf{s}_{t+j} + \sum_{j=0}^{k-1} \mathbf{v}_{r,t+1+j} \equiv \tilde{\mathbf{r}}_{t+k} \quad (71)$$

$$\mathbf{s}_{t+k} = \left(\sum_{i=0}^{k-1} \Phi_{1,s}^i \right) \cdot \Phi_{0,s} + \Phi_{1,s}^k \mathbf{s}_t + \sum_{i=0}^{k-1} \Phi_{1,s}^{k-1-i} \mathbf{v}_{s,t+1+i} \equiv \tilde{\mathbf{s}}_{t+k} \quad (72)$$

From this representation it is immediate that the state variables follow an aggregated VAR^k(1) with the following intercepts and slopes:

$$\tilde{\Phi}_{0,s} = \Phi_{0,s} \cdot \sum_{i=0}^{k-1} \Phi_{1,s}^i = \Phi_{0,s} \cdot \Theta(k) \quad \tilde{\Phi}_{1,s} = \Phi_{1,s}^k$$

where we have defined a sequence of matrices parameterized by the aggregation horizon - k , $\Theta(k) = (\mathbf{I} - \Phi_{1,s}^k) \cdot (\mathbf{I} - \Phi_{1,s})^{-1}$. The covariance matrix of the k-period shocks to the state variables is given by:

$$Var_t \left[\sum_{i=0}^{k-1} \Phi_{1,s}^{k-1-i} \mathbf{v}_{s,t+1+i} \right] = \sum_{i=0}^{k-1} \left(\Phi_{1,s}^{k-1-i} \cdot \Sigma_{ss} \cdot (\Phi_{1,s}^{k-1-i})' \right) \equiv \tilde{\Sigma}_{ss} \quad (73)$$

With these results we can re-express the k-period return as follows:

$$\begin{aligned}\tilde{\mathbf{r}}_{t+k} &= k \cdot \Phi_{0,r} + \Phi_{1,r} \cdot \sum_{j=0}^{k-1} \left(\Theta(j) \cdot \Phi_{0,s} + \Phi_{1,s}^j \cdot \mathbf{s}_t + \sum_{i=0}^{j-1} \Phi_{1,s}^{j-1-i} \mathbf{v}_{s,t+1+i} \right) + \sum_{j=0}^{k-1} \mathbf{v}_{r,t+1+j} \\ &= k \cdot \Phi_{0,r} + \Phi_{1,r} \cdot \left((k\mathbf{I} - \Theta(k)) \cdot (\mathbf{I} - \Phi_{1,s})^{-1} \cdot \Phi_{0,s} + \Theta(k) \cdot \mathbf{s}_t + \sum_{j=0}^{k-1} \sum_{i=0}^{j-1} \Phi_{1,s}^{j-1-i} \mathbf{v}_{s,t+1+i} \right) + \\ &\quad + \sum_{j=0}^{k-1} \mathbf{v}_{r,t+1+j}\end{aligned} \quad (74)$$

The form of the above expression confirms that the k-period returns also follow a first-order auto-regression with intercepts and slopes given by:

$$\tilde{\Phi}_{0,r} = k \cdot \Phi_{0,r} + \Phi_{1,r} \cdot (k\mathbf{I} - \Theta(k)) \cdot (\mathbf{I} - \Phi_{1,s})^{-1} \cdot \Phi_{0,s} \quad \tilde{\Phi}_{1,r} = \Phi_{1,r} \cdot \Theta(k)$$

The remaining two results that have to be established are the forms of the variance matrix of the k-period returns and the covariance matrix of k-period returns with the k-period ahead innovation to the state variables. We begin with the variance matrix of the k-period returns, $\tilde{\Sigma}_{rr}$. From the expression for the k-period return we see that there are two components to the shock, related to the shocks to the state variables and returns, respectively. We compute their second moments sequentially.

$$\begin{aligned} \text{Var}_t \left[\Phi_{1,r} \cdot \sum_{j=0}^{k-1} \sum_{i=0}^{j-1} \Phi_{1,s}^{j-1-i} \mathbf{v}_{s,t+1+i} \right] &= \text{Var}_t \left[\Phi_{1,r} \cdot \sum_{i=1}^{k-1} \left(\sum_{j=0}^{k-i-1} \Phi_{1,s}^j \right) \mathbf{v}_{s,t+i} \right] \\ &= \text{Var}_t \left[\Phi_{1,r} \cdot \sum_{i=1}^{k-1} \left((\mathbf{I} - \Phi_{1,s}^{k-i}) \cdot (\mathbf{I} - \Phi_{1,s})^{-1} \right) \mathbf{v}_{s,t+i} \right] \\ &= \sum_{i=1}^{k-1} \Phi_{1,r} \cdot \Theta(k-i) \cdot \Sigma_{ss} \cdot \Theta(k-i)' \cdot \Phi_{1,r}' \\ \text{Var}_t \left[\sum_{j=0}^{k-1} \mathbf{v}_{r,t+1+j} \right] &= k \cdot \Sigma_{rr} \end{aligned}$$

The covariance term appearing in the variance of the k-period return is:

$$\begin{aligned} \text{Cov}_t \left[\Phi_{1,r} \cdot \sum_{j=0}^{k-1} \sum_{i=0}^{j-1} \Phi_{1,s}^{j-1-i} \mathbf{v}_{s,t+1+i}, \sum_{j=0}^{k-1} \mathbf{v}_{r,t+1+j} \right] &= \\ = \text{Cov}_t \left[\Phi_{1,r} \cdot \sum_{i=1}^{k-1} \left((\mathbf{I} - \Phi_{1,s}^{k-i}) \cdot (\mathbf{I} - \Phi_{1,s})^{-1} \right) \mathbf{v}_{s,t+i}, \sum_{j=1}^k \mathbf{v}_{r,t+j} \right] \\ = \Phi_{1,r} \cdot \sum_{i=1}^{k-1} \Theta(k-i) \cdot \Sigma_{sr} \end{aligned}$$

Hence we arrive at the following formula for the variance of the k-period return:

$$\begin{aligned} \text{Var}_t \left[\sum_{i=0}^{k-1} \mathbf{r}_{t+1+i} \right] &= \sum_{i=1}^{k-1} \Phi_{1,r} \cdot \Theta(k-i) \cdot \Sigma_{ss} \cdot \Theta(k-i)' \cdot \Phi_{1,r}' + \\ &\quad + k \cdot \Sigma_{rr} + 2\Phi_{1,r} \cdot \sum_{i=1}^{k-1} \Theta(k-i) \cdot \Sigma_{sr} \equiv \tilde{\Sigma}_{rr} \end{aligned} \quad (75)$$

Lastly, we turn to the covariance of the k-period shocks to the state variables and the shocks to the k-period return:

$$\text{Cov}_t \left[\sum_{i=0}^{k-1} \Phi_{1,s}^{k-1-i} \mathbf{v}_{s,t+1+i}, \Phi_{1,r} \cdot \sum_{j=0}^{k-1} \sum_{i=0}^{j-1} \Phi_{1,s}^{j-1-i} \mathbf{v}_{s,t+1+i} + \sum_{j=0}^{k-1} \mathbf{v}_{r,t+1+j} \right] \equiv \tilde{\Sigma}_{sr}$$

Using the properties of the covariance operator, computing the sums and simplifying the notation using the

$\Theta(i)$ matrices allows us to re-express the covariance as:

$$\begin{aligned}
\tilde{\Sigma}_{sr} &= Cov_t \left[\sum_{i=1}^k \Phi_{1,s}^{k-i} \mathbf{v}_{s,t+i}, \Phi_{1,r} \cdot \sum_{i=1}^{k-1} \left((\mathbf{I} - \Phi_{1,s}^{k-i}) \cdot (\mathbf{I} - \Phi_{1,s})^{-1} \right) \mathbf{v}_{s,t+i} + \sum_{j=1}^k \mathbf{v}_{r,t+j} \right] \\
&= \sum_{i=1}^{k-1} \Phi_{1,s}^{k-i} \cdot \Sigma_{ss} \cdot \Theta(k-i)' \cdot \Phi_{1,r}' + \sum_{j=1}^k \Phi_{1,s}^{k-j} \Sigma_{sr} \\
&= \sum_{i=1}^{k-1} \Phi_{1,s}^{k-i} \cdot \Sigma_{ss} \cdot \Theta(k-i)' \cdot \Phi_{1,r}' + \Theta(k) \cdot \Sigma_{sr} \equiv \tilde{\Sigma}_{sr}
\end{aligned} \tag{76}$$

Ultimately, we are in position to conclude that the aggregated VAR^k(1) model takes the following form:

$$\begin{bmatrix} \tilde{\mathbf{r}}_{t+k} \\ \tilde{\mathbf{s}}_{t+k} \end{bmatrix} = \begin{bmatrix} \tilde{\Phi}_{0,r} \\ \tilde{\Phi}_{0,s} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{(1+n) \times (1+n)} & \tilde{\Phi}_{1,r} \\ \mathbf{0}_{m \times (1+n)} & \tilde{\Phi}_{1,s} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{0}_{(1+n) \times 1} \\ \mathbf{s}_t \end{bmatrix} + \tilde{\mathbf{v}}_{t+k}$$

where the covariance matrix of the shocks to the aggregated system is given by:

$$\tilde{\Sigma}_v = \begin{bmatrix} \tilde{\Sigma}_{rr} & \tilde{\Sigma}_{rs} \\ \tilde{\Sigma}'_{rs} & \tilde{\Sigma}_{ss} \end{bmatrix}$$

Table 1: **Sources of time-series used in the construction of state variables and asset returns for the VAR(1) model.** The first panel provides the origin of the raw data used to estimate the model, while the second describes the variables included in the VAR(1) along with the acronyms used in tables reporting model estimates.

Time series	Source	Series name
Inflation	CRSP	<i>cpiret</i>
Returns on the 30-day T-bill	CRSP	<i>t30ret</i>
Returns on the 5-yr Treasury bond	CRSP	<i>b5ret</i>
30-day T-bill yield (in percentage points)	CRSP	<i>tb30yield</i>
5-year discount bond yield (in percentage points)	CRSP	<i>yield5</i>
Returns on the value-weighted CRSP Index	constructed from FF25 data	-
Returns on the value stock index	constructed from FF25 data	-
Returns on the growth stock index	constructed from FF25 data	-
S&P 500 PE ratio (P/10-yr E)	Shiller website	-

Variable definition	VAR(1) Acronym
Log ex-post real return on the 30-day T-bill	t30_realret
Log excess return on the value-weighted CRSP index	M-rf
Log excess return on the value stock index	V-rf
Log excess return on the growth stock index	G-rf
Log excess return on the 5-yr Treasury bond	B5-rf
Log value/growth return spread	V-G
Log ex-post real return on the growth stock index	G real

The return series for the value and growth indices are constructed as a value-weighted average of the returns on the six book-to-market/size sorted Fama-French portfolios. The weights are computed as the ratio of the market capitalizations of the underlying portfolios to the market capitalization of the respective index (defined in the caption to Figure 1). Hence the return on the value and growth indices at time t is defined as follows:

$$\begin{aligned}
 r_t^V &= r_t^{SMHI} \left(\frac{ME_t^{SMHI}}{ME_t^V} \right) + \frac{1}{2} r_t^{SMME} \left(\frac{ME_t^{SMME}}{ME_t^V} \right) + r_t^{BIHI} \left(\frac{ME_t^{BIHI}}{ME_t^V} \right) + \frac{1}{2} r_t^{BIME} \left(\frac{ME_t^{BIME}}{ME_t^V} \right) \\
 r_t^G &= r_t^{SMLO} \left(\frac{ME_t^{SMLO}}{ME_t^G} \right) + \frac{1}{2} r_t^{SMME} \left(\frac{ME_t^{SMME}}{ME_t^G} \right) + r_t^{BILO} \left(\frac{ME_t^{BILO}}{ME_t^G} \right) + \frac{1}{2} r_t^{BIME} \left(\frac{ME_t^{BIME}}{ME_t^G} \right)
 \end{aligned}$$

Table 2: **Summary statistics for variables used in VAR estimation.** The dataset is comprised of monthly data from 1952:12 to 2003:12 (613 data points).

Assets	Mean (% p.a.)	St. dev. (% p.a.)
Log (nominal) return on 30-day T-bill	5.07%	0.83%
Log (ex-post real) return on 30-day T-bill	1.28%	0.98%
Log return on market portfolio	10.68%	14.99%
Log return on growth portfolio	10.36%	15.50%
Log return on value portfolio	12.81%	14.20%
Log value-growth spread return	2.45%	6.74%
Log return on 5-year bond	6.46%	5.23%
State variables	Mean (% p.a.)	St. dev. (% p.a.)
Yield spread (5-yr Bond – 30-day T-bill)	1.31%	1.09%
Log PE ratio	2.82	0.41
Nominal yield (30-day T-bill)	4.98%	2.77%
Other	Mean (% p.a.)	St. dev. (% p.a.)
Log CPI	3.78%	1.15%
Growth share (in MKT portfolio)	70.18%	7.68%
Value share (in MKT portfolio)	29.82%	7.68%

	Sharpe Ratio
Market portfolio (MKT)	0.449
Growth portfolio (G)	0.419
Value portfolio (V)	0.616
5-yr Bond (B5)	0.295

Table 3: **Estimates of the equity-only VAR(1) system.** The investable universe only includes the value (V) and growth (G) stock portfolios. The VAR estimates are obtained using monthly data from 1952:12 to 2003:12 (613 data points). The top panel contains the estimates of the VAR coefficients with t-statistics (round brackets) and bootstrap estimates of the 95%-confidence intervals (square brackets). The t-statistics are computed using Newey-West standard errors with 12 lags. The bootstrap estimates are produced from 10,000 paths (with length equal to the historical sample) simulated under the assumption that the estimated VAR process is the true data generating process. The rightmost column reports the regression R^2 and the p-value of the F-test for no predictability in parentheses. The bottom panel contains the residual correlation matrix for the VAR(1) shocks; entries on the diagonal correspond to monthly residual standard deviations.

A. VAR Coefficient Estimates							
	G real	V-G	YSPR	PE	t30_yield	t30_realret	R^2
G real	0.0201 (0.45)	-0.1319 (-1.33)	0.0021 (1.10)	-0.0142 (-2.24)	-0.2909 (-3.51)	1.8668 (3.39)	4.81% (0.00)
V-G	[-0.0582, 0.0875] 0.0265 (1.17)	[-0.3000, 0.0368] 0.1434 (2.86)	[-0.0013, 0.0055] -0.0006 (-0.77)	[-0.0344, -0.0091] 0.0004 (0.08)	[-0.5064, -0.1167] 0.0422 (0.96)	[0.7415, 2.9656] -0.3958 (-1.29)	2.46% (0.01)
YSPR	[-0.0053, 0.0575] -0.3900 (-0.64)	[0.0625, 0.2067] 0.6553 (0.59)	[-0.0021, 0.0008] 0.8656 (34.22)	[-0.0036, 0.0064] -0.0678 (-1.20)	[-0.0260, 0.1291] 0.0160 (0.02)	[-0.8731, 0.0906] 3.3763 (0.54)	75.46% (0.00)
PE	[-1.2569, 0.5124] 0.4166 (16.39)	[-1.3595, 2.7278] -0.0448 (-0.69)	[0.8120, 0.8949] 0.0024 (2.15)	[-0.2693, 0.0282] 0.9920 (264.64)	[-2.1041, 2.5877] -0.1407 (-2.81)	[-10.3248, 16.9942] 1.0233 (2.73)	99.53% (0.00)
t30_yield	[0.3670, 0.4609] 0.0157 (1.87)	[-0.1560, 0.0597] -0.0028 (-0.19)	[0.0001, 0.0045] 0.0010 (3.02)	[0.9789, 0.9951] 0.0005 (0.87)	[-0.2775, -0.0296] 0.9949 (79.78)	[0.3029, 1.7246] -0.1789 (-2.08)	95.13% (0.00)
t30_realret	[0.0053, 0.0256] 0.0036 (1.17)	[-0.0259, 0.0200] -0.0005 (-0.08)	[0.0005, 0.0015] 0.0003 (2.33)	[-0.0007, 0.0029] 0.0010 (2.39)	[0.9539, 1.0112] 0.0262 (3.94)	[-0.3296, -0.0240] 0.2758 (4.59)	15.81% (0.00)
	[-0.0006, 0.0080]	[-0.0104, 0.0095]	[0.0002, 0.0006]	[0.0001, 0.0015]	[0.0175, 0.0387]	[0.2021, 0.3310]	

B. VAR Residual Correlations and Standard Deviations						
	G real	V-G	YSPR	PE	t30_yield	t30_realret
G real	0.0442	-0.3927	-0.0618	0.7732	-0.0240	0.1544
V-G	-	0.0192	0.0384	-0.1911	-0.0550	-0.0021
YSPR	-	-	0.5418	-0.0268	-0.8149	0.0969
PE	-	-	-	0.0283	-0.0247	0.1303
t30_yield	-	-	-	-	0.0061	-0.1391
t30_realret	-	-	-	-	-	0.0026

Table 5: **Composition of the one-period tangency and global minimum variance (GMV) portfolios.**

Investable universe	Asset	Tangency	GMV
(V, G)	V	572.23%	90.35%
	G	-472.23%	9.65%
(T-bill, V, G, B5)	T-bill	-726.84%	101.76%
	V	730.51%	-0.51%
	G	-434.14%	-0.12%
	B5	530.47%	-1.13%

Table 6: **Welfare simulation results for the equity-only case by horizon, risk-aversion and rebalancing frequency.** Panel A reports the certainty equivalents of wealth for the dynamic strategy with monthly rebalancing, which serves as the base case for the welfare loss computations. Panels B1-B4 report the welfare losses of dynamic strategies with lower rebalancing frequencies and the constant proportion strategy. The welfare loss (η) is computed as the percentage loss in the certainty equivalent of wealth between the proposed rebalancing scheme and the base case. Panels C1-C4 compute the maximal monthly management fee (ϕ) an agent would be willing to pay in order to gain access to the dynamic monthly strategy. The fee is computed as $\phi = 1 - (1 + \eta)^{\frac{1}{\gamma}}$. The results are obtained by Monte Carlo simulation using 20,000 VAR paths sampled using the method of antithetic variates, with the same path set being used to evaluate the welfare loss in all cases. The certainty equivalent of wealth is computed by evaluating the mean utility realized across the simulated paths and computing, $W_{CE} = u^{-1}(E[u(\bar{W}_T)])$. γ denotes the investor's coefficient of relative risk aversion.

	γ	1mo	6mo	12mo	24mo	60mo	120mo	300mo
A. Certainty Equivalent								
A. Dynamic / monthly	1	2.744	2.878	3.054	3.451	4.982	9.548	81.002
	2	1.013	1.080	1.168	1.364	2.154	4.567	41.200
	5	1.005	1.027	1.055	1.104	1.264	1.580	3.577
	10	0.999	0.990	0.978	0.937	0.841	0.810	1.160
B. Welfare loss								
B1. Dynamic / quarterly	1	-	-0.7%	-1.4%	-3.5%	-8.5%	-16.0%	-34.3%
	2	-	-0.3%	-0.6%	-1.7%	-4.4%	-8.7%	-19.8%
	5	-	-0.1%	-0.2%	-0.5%	-2.2%	-4.2%	-9.7%
	10	-	-0.0%	-0.2%	0.1%	-2.0%	-2.6%	-3.3%
B2. Dynamic / semi-annual	1	-	-1.3%	-2.6%	-5.4%	-13.0%	-23.3%	-47.3%
	2	-	-0.6%	-1.2%	-2.6%	-6.8%	-12.9%	-27.8%
	5	-	-0.2%	-0.5%	-0.7%	-3.0%	-6.8%	-14.1%
	10	-	-0.0%	-0.3%	0.5%	-1.1%	-5.5%	-9.3%
B3. Dynamic / annual	1	-	-	-3.2%	-7.0%	-15.8%	-28.5%	-56.0%
	2	-	-	-1.6%	-3.5%	-8.5%	-16.1%	-34.7%
	5	-	-	-0.7%	-1.0%	-4.1%	-9.1%	-18.5%
	10	-	-	-0.6%	0.1%	-3.5%	-5.6%	-11.9%
B4. Constant proportion	1	0.0%	-0.5%	-1.6%	-4.9%	-14.9%	-29.9%	-61.9%
	2	0.0%	-0.2%	-0.8%	-2.5%	-7.9%	-16.9%	-39.1%
	5	0.0%	-0.1%	-0.3%	-1.0%	-3.4%	-8.1%	-21.7%
	10	0.0%	-0.0%	-0.2%	-0.6%	-1.7%	-2.8%	-15.8%
C. Maximal monthly fee								
C1. Dynamic / quarterly	1	-	0.11%	0.12%	0.15%	0.15%	0.15%	0.14%
	2	-	0.05%	0.05%	0.07%	0.07%	0.08%	0.07%
	5	-	0.01%	0.02%	0.02%	0.04%	0.04%	0.03%
	10	-	0.00%	0.01%	0.00%	0.03%	0.02%	0.01%
C2. Dynamic / semi-annual	1	-	0.22%	0.22%	0.23%	0.23%	0.22%	0.21%
	2	-	0.10%	0.10%	0.11%	0.12%	0.11%	0.11%
	5	-	0.03%	0.04%	0.03%	0.05%	0.06%	0.05%
	10	-	0.00%	0.03%	0.00%	0.02%	0.05%	0.03%
C3. Dynamic / annual	1	-	-	0.27%	0.30%	0.29%	0.28%	0.27%
	2	-	-	0.14%	0.15%	0.15%	0.15%	0.14%
	5	-	-	0.06%	0.04%	0.07%	0.08%	0.07%
	10	-	-	0.05%	0.00%	0.06%	0.05%	0.04%
C4. Constant proportion	1	0.00%	0.08%	0.14%	0.21%	0.27%	0.30%	0.32%
	2	0.00%	0.04%	0.07%	0.11%	0.14%	0.15%	0.17%
	5	0.00%	0.01%	0.03%	0.04%	0.06%	0.07%	0.08%
	10	0.00%	0.01%	0.01%	0.02%	0.03%	0.02%	0.06%

Table 7: **Time-series variance of the portfolio allocation to value (V)**. The table presents a decomposition of the time-series variance of the portfolio allocation to value (V) in the equity-only case for an investor with relative risk aversion of γ who chooses her portfolio as if there perpetually were τ periods remaining to the investment horizon. The table reports the total variance realized over the period 1952:12 to 2003:12 for which VAR data is available, and the fraction of the portfolio weight variance due to the intertemporal hedging component. Since the growth (G) portfolio weight for any horizon is defined as the residual between unity and the value portfolio weight, the analogous variances for the growth portfolio weight are identical.

γ	Variable	12mo	60mo	120mo	300mo
2	$\sigma(\alpha_V^{tot})$	406.52%	406.62%	406.62%	406.60%
	$\frac{\sigma(\alpha_V^h)}{\sigma(\alpha_V^{tot})}$	0.0237	0.0231	0.0231	0.0232
10	$\sigma(\alpha_V^{tot})$	80.86%	80.89%	80.89%	80.89%
	$\frac{\sigma(\alpha_V^h)}{\sigma(\alpha_V^{tot})}$	0.0383	0.0374	0.0375	0.0376
500	$\sigma(\alpha_V^{tot})$	1.62%	1.62%	1.62%	1.62%
	$\frac{\sigma(\alpha_V^h)}{\sigma(\alpha_V^{tot})}$	0.0414	0.0405	0.0405	0.0406

Table 8: **Welfare simulation results for the T-bills, bonds, and equities case by horizon, risk-aversion and rebalancing frequency.** Panel A reports the certainty equivalents of wealth for the dynamic strategy with monthly rebalancing, which serves as the base case for the welfare loss computations. Panels B1-B4 report the welfare losses of dynamic strategies with lower rebalancing frequencies and the constant proportion strategy. The welfare loss (η) is computed as the percentage loss in the certainty equivalent of wealth between the proposed rebalancing scheme and the base case. Panels C1-C4 compute the maximal monthly management fee (ϕ) an agent would be willing to pay in order to gain access to the dynamic monthly strategy. The fee is computed as $\phi = 1 - (1 + \eta)^{\frac{1}{\gamma}}$. The results are obtained by Monte Carlo simulation using 20,000 VAR paths sampled using the method of antithetic variates, with the same path set being used to evaluate the welfare loss in all cases. The certainty equivalent of wealth is computed by evaluating the mean utility realized across the simulated paths and computing, $W_{CE} = u^{-1}(E[u(\bar{W}_T)])$. γ denotes the investor's coefficient of relative risk aversion

	γ	1mo	6mo	12mo	24mo	60mo	120mo	300mo
A. Certainty Equivalent								
A. Dynamic / monthly	1	2.831	3.478	4.399	7.133	31.178	368.769	8.608×10^5
	2	1.027	1.173	1.372	1.878	4.861	24.759	4481.890
	5	1.011	1.065	1.135	1.288	1.918	3.916	53.223
	10	1.006	1.035	1.071	1.148	1.436	2.136	10.590
B. Welfare loss								
B1. Dynamic / quarterly	1	-	-3.6%	-6.1%	-11.5%	-26.0%	-44.3%	-77.0%
	2	-	-1.6%	-3.1%	-5.7%	-13.1%	-25.6%	-52.8%
	5	-	-0.6%	-1.2%	-2.1%	-5.1%	-11.5%	-28.9%
	10	-	-0.3%	-0.6%	-1.0%	-2.7%	-6.4%	-18.9%
B2. Dynamic / semi-annual	1	-	-6.4%	-11.3%	-21.0%	-44.2%	-68.8%	-94.6%
	2	-	-2.9%	-5.7%	-10.9%	-24.3%	-44.7%	-79.3%
	5	-	-1.1%	-2.3%	-4.5%	-10.6%	-22.0%	-54.6%
	10	-	-0.6%	-1.2%	-2.2%	-5.7%	-11.7%	-38.0%
B3. Dynamic / annual	1	-	-	-16.7%	-30.6%	-59.8%	-83.3%	-98.9%
	2	-	-	-8.7%	-16.5%	-36.0%	-60.1%	-91.4%
	5	-	-	-3.6%	-7.1%	-16.8%	-33.3%	-72.9%
	10	-	-	-1.9%	-3.6%	-9.2%	-19.2%	-55.1%
B4. Constant proportion	1	0.0%	-1.2%	-5.6%	-20.2%	-60.9%	-90.5%	-99.9%
	2	0.0%	-0.5%	-3.0%	-10.9%	-36.2%	-69.8%	-98.2%
	5	0.0%	-0.2%	-1.3%	-4.6%	-16.7%	-39.9%	-88.3%
	10	0.0%	-0.1%	-0.6%	-2.3%	-9.4%	-23.0%	-73.2%
C. Maximal monthly fee								
C1. Dynamic / quarterly	1	-	0.61%	0.53%	0.51%	0.50%	0.49%	0.49%
	2	-	0.27%	0.26%	0.24%	0.23%	0.25%	0.25%
	5	-	0.10%	0.10%	0.09%	0.09%	0.10%	0.11%
	10	-	0.05%	0.05%	0.04%	0.04%	0.06%	0.07%
C2. Dynamic / semi-annual	1	-	1.09%	1.00%	0.98%	0.97%	0.97%	0.97%
	2	-	0.50%	0.49%	0.48%	0.46%	0.49%	0.52%
	5	-	0.19%	0.20%	0.19%	0.19%	0.21%	0.26%
	10	-	0.10%	0.10%	0.09%	0.10%	0.10%	0.16%
C3. Dynamic / annual	1	-	-	1.51%	1.51%	1.51%	1.48%	1.49%
	2	-	-	0.75%	0.75%	0.74%	0.76%	0.81%
	5	-	-	0.31%	0.31%	0.31%	0.34%	0.43%
	10	-	-	0.16%	0.15%	0.16%	0.18%	0.27%
C4. Constant proportion	1	0.00%	0.20%	0.48%	0.94%	1.55%	1.95%	2.44%
	2	0.00%	0.09%	0.25%	0.48%	0.75%	0.99%	1.35%
	5	0.00%	0.04%	0.11%	0.20%	0.30%	0.42%	0.71%
	10	0.00%	0.02%	0.05%	0.10%	0.16%	0.22%	0.44%

Figure 1: **Term-structure of the annualized log population arithmetic and geometric mean excess returns on the aggregate stock market (Mkt) and the 5-yr Treasury bond (B5).**

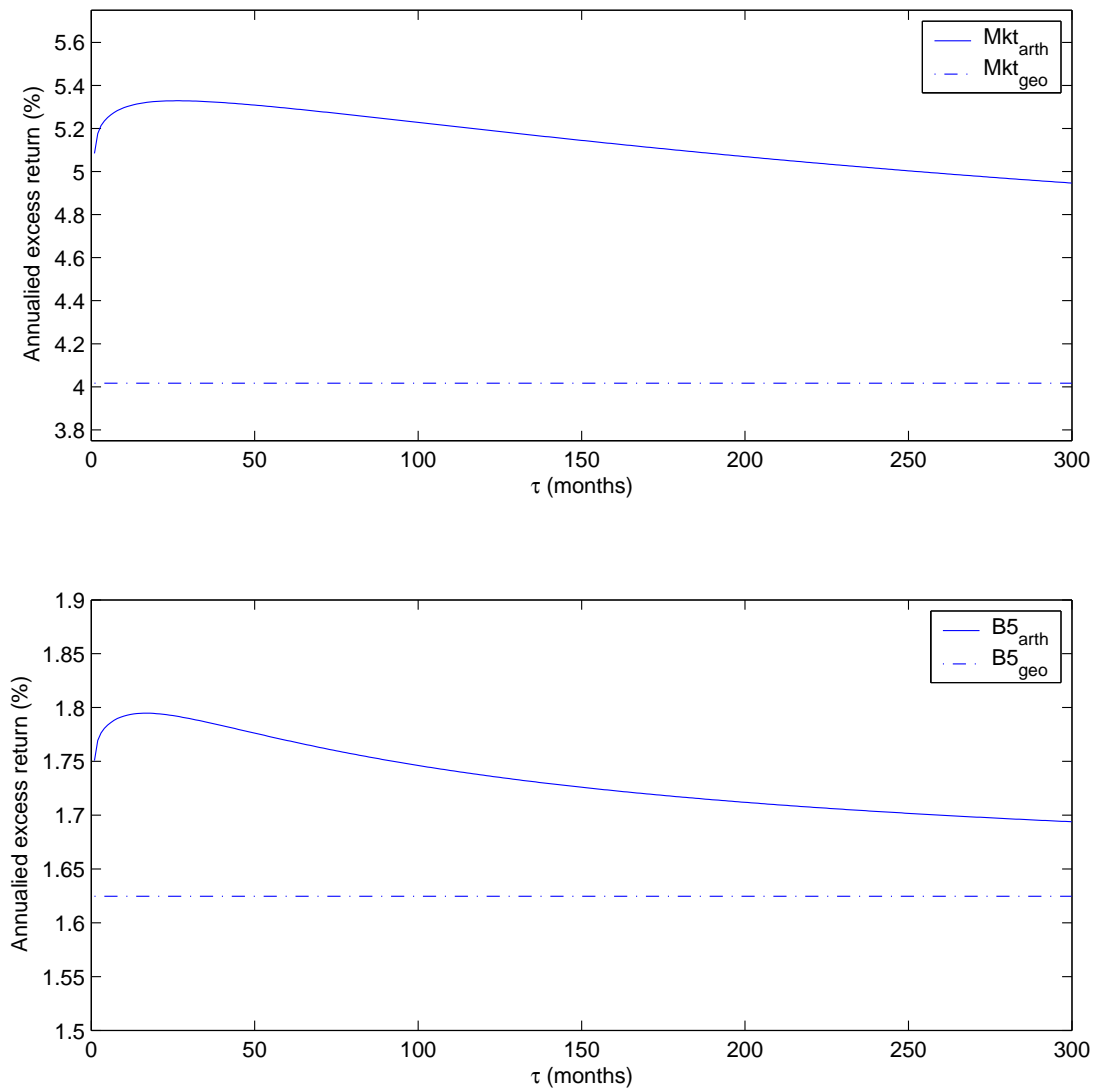


Figure 2: **Market capitalization of value and growth stock portfolios as a fraction of total market capitalization.** The value (V) and growth (G) indices used in the VAR estimation are constructed using the data available on Prof. Ken French's website. We use the time series of data available for six portfolios resulting from an intersection of two portfolios formed on size (market equity) and three portfolios formed on the ratio of book equity to market equity. The *value index* is defined to include the SMHI and BIHI portfolios, as well as, one half of the SMME and BIME portfolios. The *growth index* has the complementary composition and includes the SMLO and BILO portfolios, and one half of the SMME and BIME portfolios. To construct the time series of the market capitalizations of each of the indices we use the data on number of companies in each of the sub-portfolios (N_t) and their average market capitalization (\overline{ME}_t). Thus the market capitalization of the value index at time t is defined as follows: $ME_t^V = N_t^{SMHI} \cdot \overline{ME}_t^{SMHI} + \frac{1}{2} N_t^{SMME} \cdot \overline{ME}_t^{SMME} + N_t^{BIHI} \cdot \overline{ME}_t^{BIHI} + \frac{1}{2} N_t^{BIME} \cdot \overline{ME}_t^{BIME}$. The market capitalization of the growth index (G) has the complementary definition.

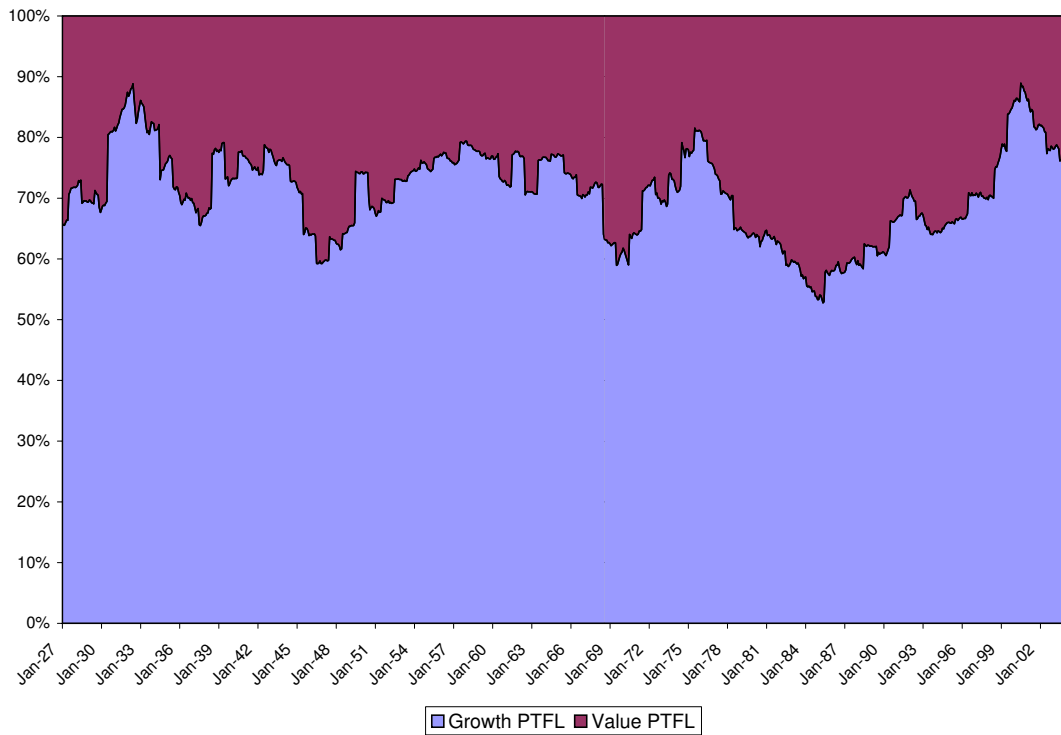


Figure 3: Mean percentage optimal portfolio allocations in the equity-only case.

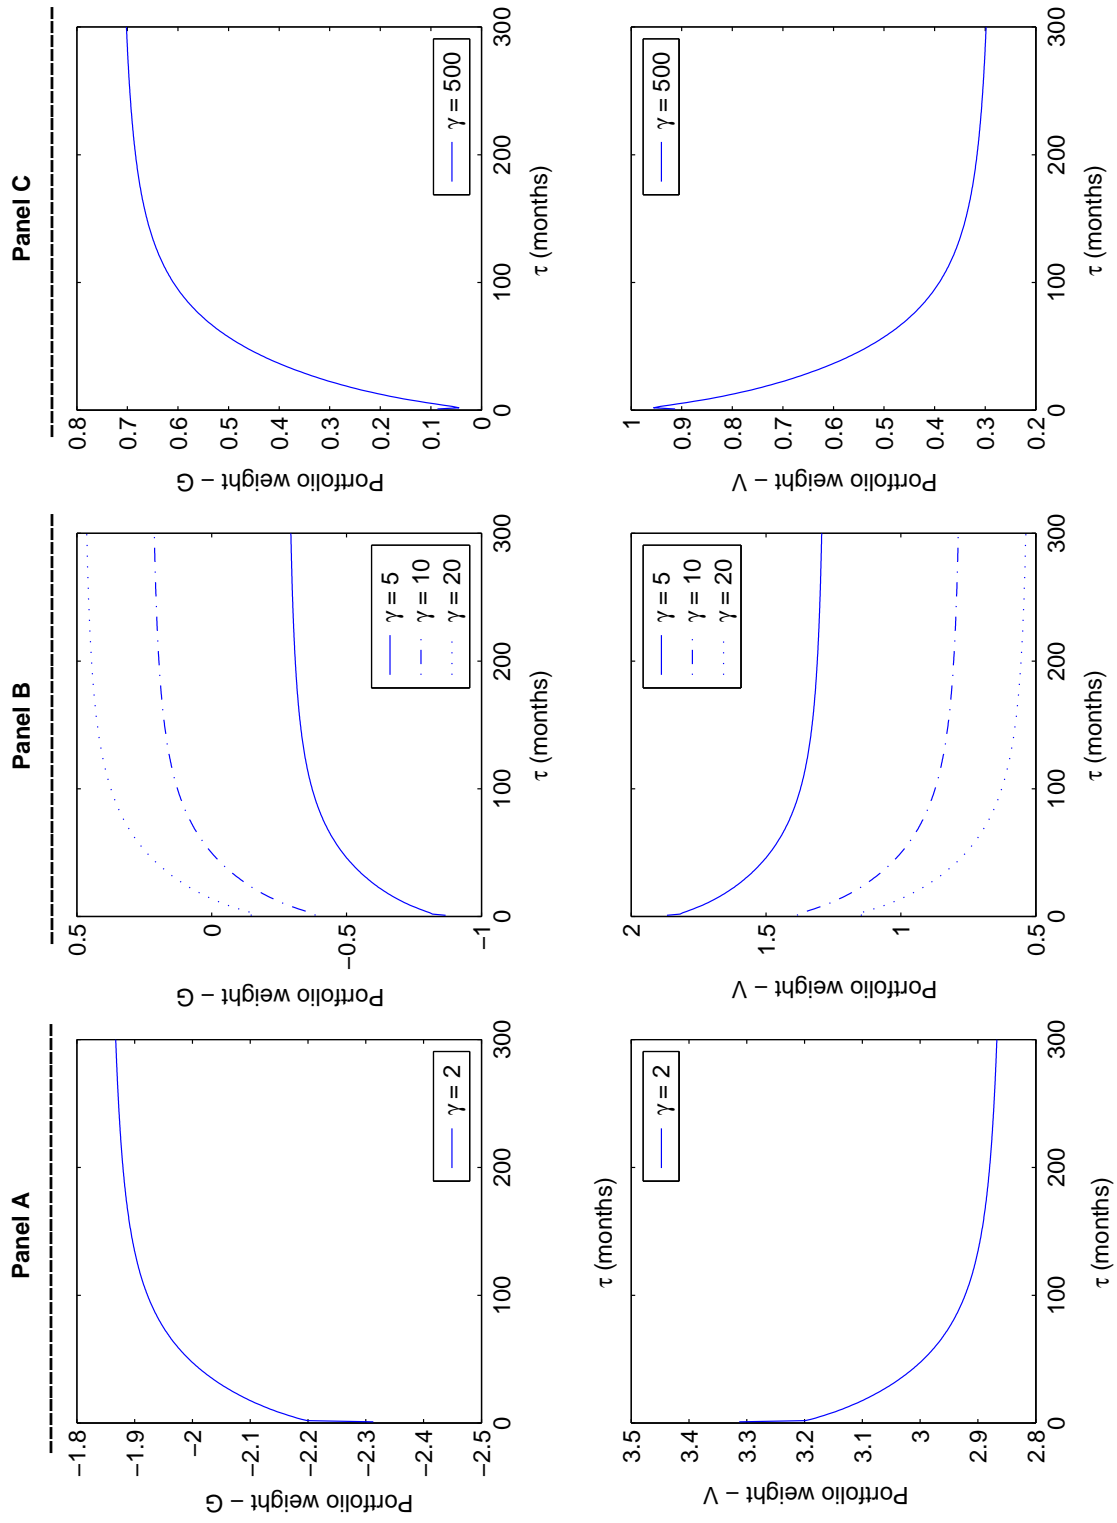


Figure 4: Mean percentage optimal allocations to cash, bonds, growth and value for investors with $\gamma = 2$ (Panel A), and implied mean percentage optimal allocation to the aggregate stock market and value tilt (Panel B).

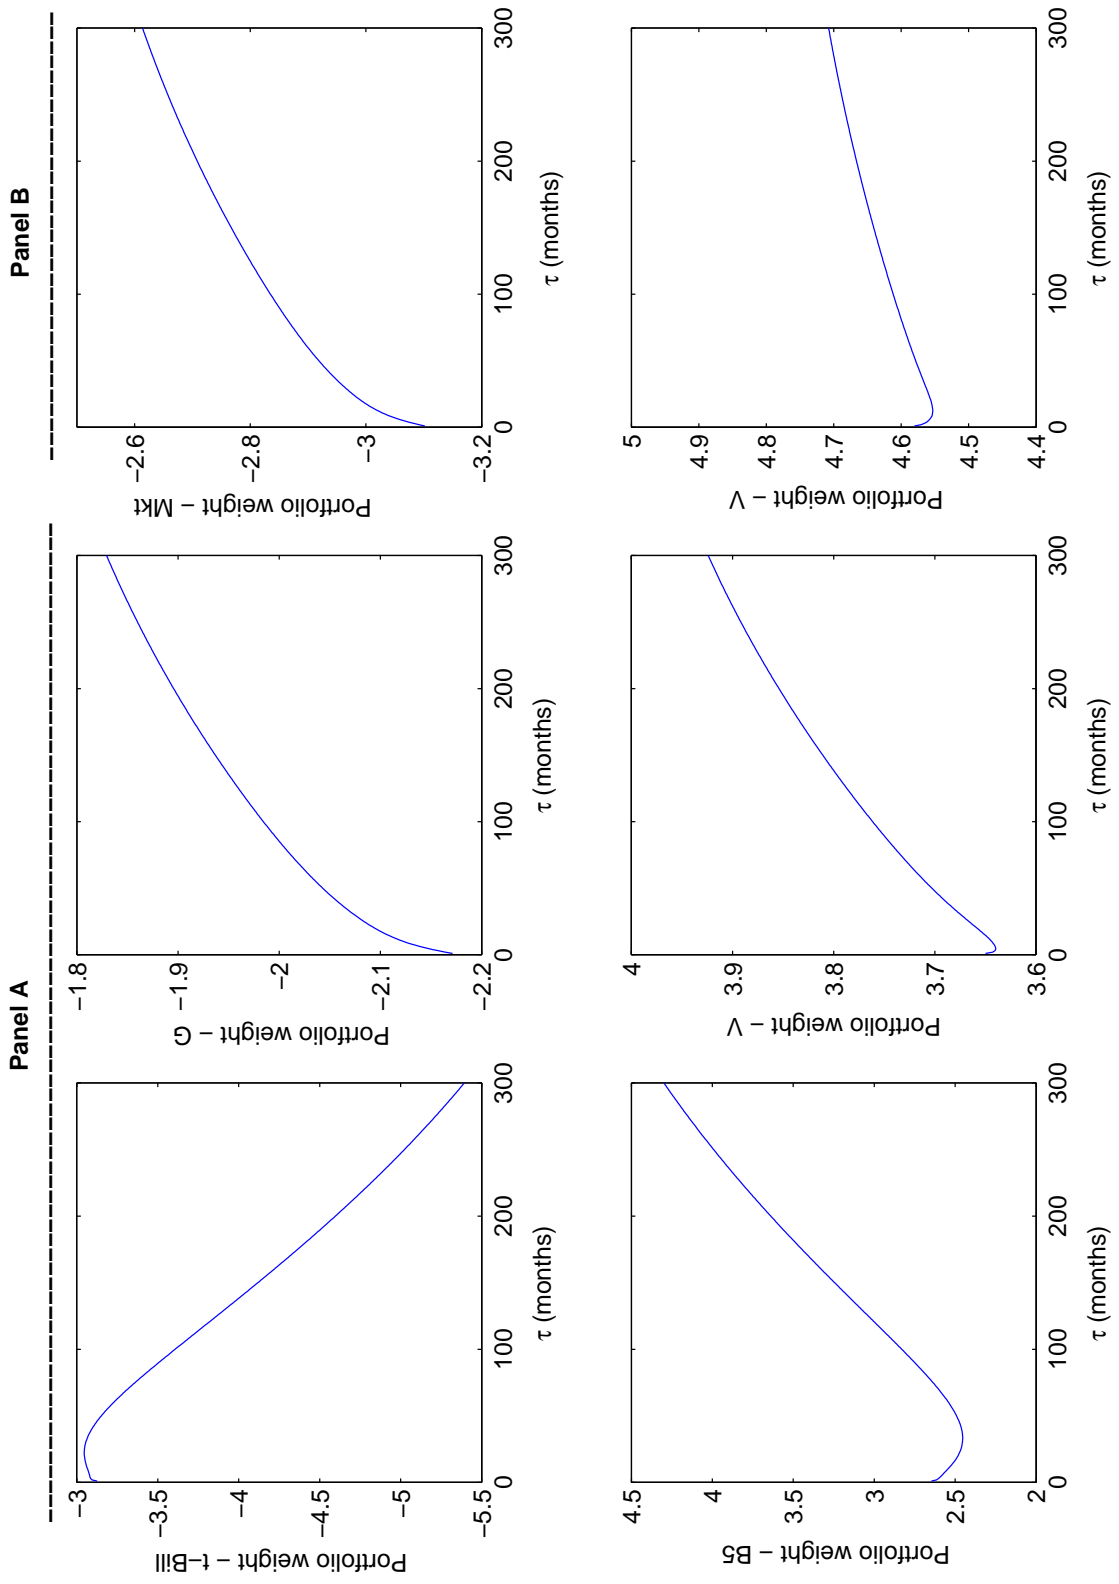


Figure 5: Mean percentage optimal allocations to cash, bonds, growth and value for investors with $\gamma = \{5, 10, 20\}$ (Panel A), and implied mean percentage optimal allocation to the aggregate stock market and value tilt (Panel B).

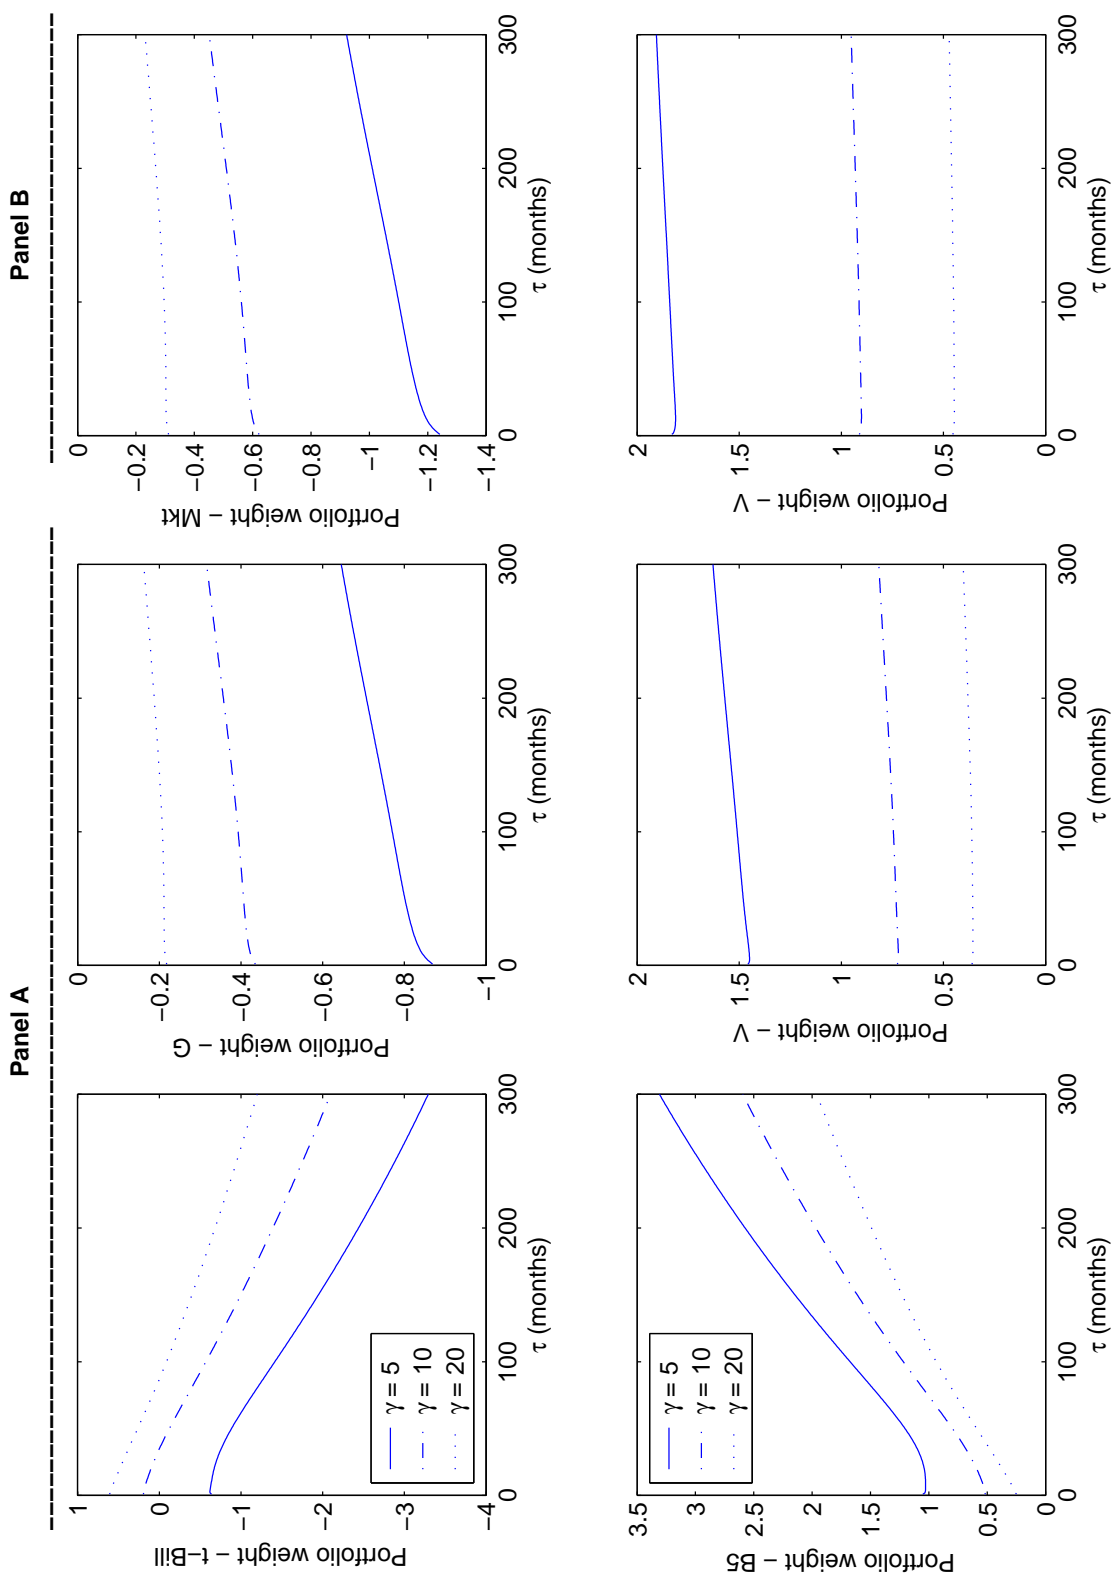


Figure 6: Mean percentage allocations to cash, bonds, the aggregate stock market, and the optimal value tilt for investors with $\gamma = 500$.

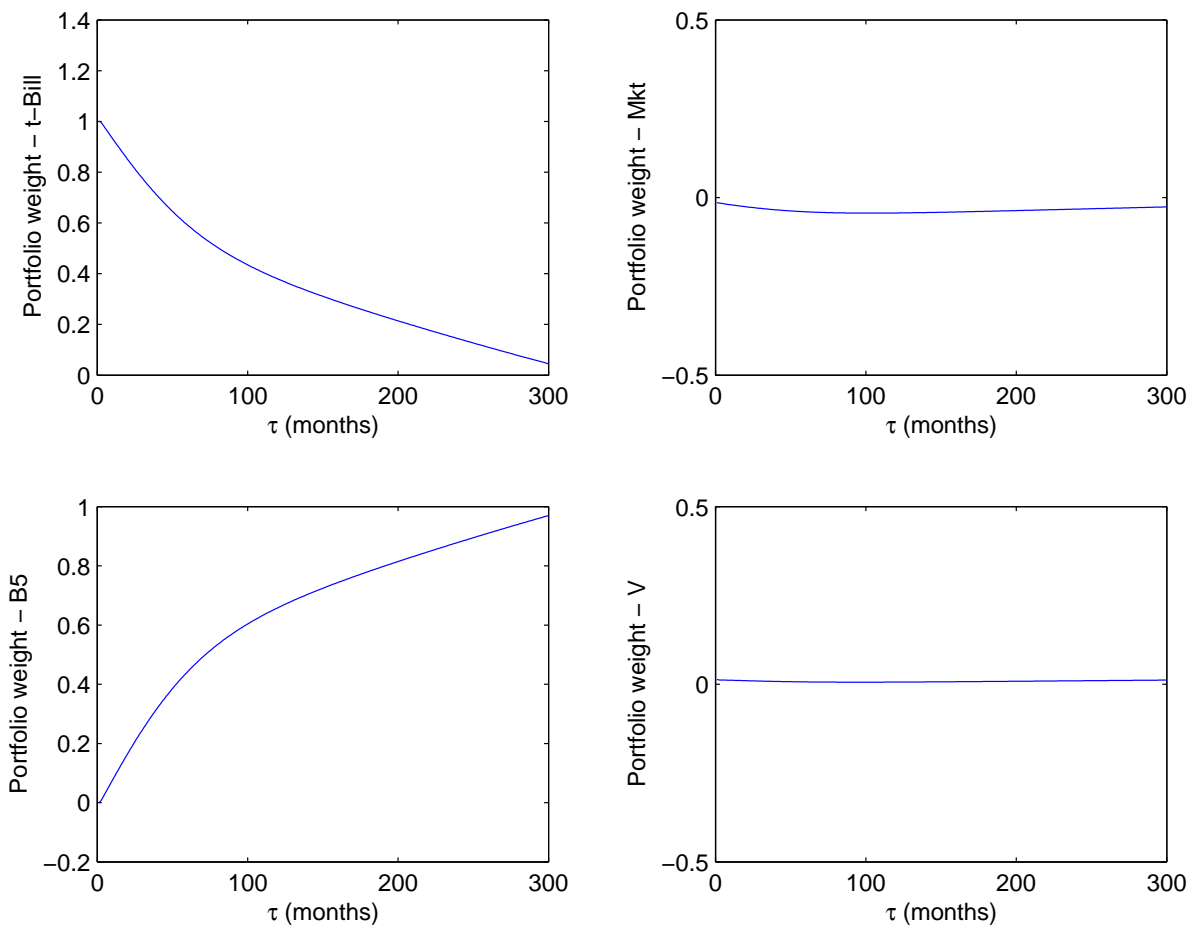


Figure 7: Comparison of mean portfolio allocations under the constant proportion strategy (dashed line) and the optimal dynamic strategy (solid line) for investors with $\gamma = 2$ (Panel A) and $\gamma = 500$ (Panel B) in the equity-only case.

