

First Draft: March 1999

This Draft: February 2000

**Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands
Induced by Return Predictability**

Anthony W. Lynch*

*New York University, 44 West 4th St Suite 9-190, New York NY 10012; ph: (212) 998-0350;
email: alynch@stern.nyu.edu

The first draft of paper was written while the author was visiting Columbia University. I would like to thank Doron Avramov, Michael Brandt, Ned Elton, Wayne Ferson, Matt Richardson, Rob Stambaugh, Dimitri Vayanos, Jessica Wachter, and seminar participants at Carnegie-Mellon, the Wharton School, the 2000 AFA Meetings, the November 1999 NBER Asset Pricing Conference, the NYU Friday Seminar Series, the NYU Friday Morning Macro/Finance Seminar, the Federal Reserve of New York and the GSAM Quantitative Research Group Seminar for helpful comments.

Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands induced by Return Predictability

Abstract

This paper examines portfolio allocation across equity portfolios formed on the basis of characteristics like size and book-to-market. In particular, the paper assesses the impact of return predictability on portfolio choice for a multi-period investor with a coefficient of relative risk aversion of 4. Compared to the investor's allocation in her last period, return predictability with dividend yield causes the investor early in life to tilt her risky-asset portfolio away from high book-to-market stocks and away from small stocks. These results are explained using Merton's (1973) characterization of portfolio allocation by a multiperiod investor in a continuous time setting. Abnormal returns relative to the investor's optimal early-life portfolio are also calculated. These abnormal returns are found to exhibit the same cross-sectional patterns as abnormal returns calculated relative to the market portfolio: higher for small than large firms, and higher for high than low book-to-market firms. Thus, hedging demand may be a partial explanation for the high expected returns documented empirically for small firms and high book-to-market firms. However, even with this hedging demand, the investor wants to short-sell the low book-to market portfolio to hold the high book-to-market portfolio. The utility costs of using a value-weighted equity index or of ignoring predictability are also calculated. An investor using a value-weighted equity index would give up a much larger fraction of her wealth to have access to book-to-market portfolios than size portfolios. Finally, an investor would give up a much larger fraction of her wealth to have access to dividend yield information than term spread information.

JEL classification: G11; G12.

Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands induced by Return Predictability

1 Introduction

Fama and French (1992) find that size and book-to-market explain cross-sectional variation in expected return over and above CAPM beta. One explanation for this result is that investors care about more than mean and standard deviation when choosing their portfolio. Merton (1973) and Fama (1970) describe conditions under which investors also care about the covariance of their portfolios with a set of state variables. If returns are predictable using a set of lagged instruments and investors are multi-period optimizers, then investors care about the covariance of their portfolios with those instruments. Empirical research indicates that U.S. equity returns are predictable and that investor horizons are longer than a month.¹ However, for these multi-period considerations to affect expected returns in equilibrium, they must have a large impact on portfolio choice. This paper examines how an investor's multi-period horizon affects portfolio choice when returns are predictable and calibrated to U.S. data. The calibrated equity portfolios are chosen to exhibit cross-sectional variation in firm size and firm book-to market. The investor has constant relative risk aversion utility with a coefficient of 4.

Several recent papers have considered portfolio allocations by a multiperiod investor confronted with return predictability that is calibrated to U.S. data. However, in all these papers, the only equity portfolio available to the investor is the market portfolio. My paper is the first to consider portfolio allocation by a multiperiod investor who has access to more than one domestic equity portfolio. Consequently, hedging demand can affect not just the amount allocated to equities by the investor but the composition of the equity portfolio as well. Since expected returns have been found empirically to vary with size and book-to-market, the investor is given access to portfolios formed on the basis of these two firm characteristics. It follows that the paper can potentially say something about hedging demand as an explanation for the size and book-to-market effects. Thus, this paper is the first portfolio choice paper to consider whether the size and book-to-market effects can be explained by the hedging demand induced by predictability.

¹Campbell (1987) and Fama and French (1989), among others, find that stock return variation can be explained by the one-month Treasury bill rate, the term premium, and the dividend yield.

The paper has several interesting results. Compared to the investor's allocation in her last period, return predictability with dividend yield causes the investor early in life to tilt her risky-asset portfolio away from high book-to-market stocks and away from small stocks. Further, while the investor's optimal last-period portfolio is indistinguishable from the conditional minimum-variance frontier in conditional mean-standard deviation space, her optimal early-in-life portfolio has a larger negative covariance with dividend yield than the conditional minimum-variance portfolio with the same expected return. Treating the young investor as the representative agent, CAPM abnormal returns are calculated using the investor's optimal early-life portfolio as the market portfolio. Since hedging demand produces an optimal early-in-life portfolio that lies inside the conditional mean-variance frontier, these abnormal returns must be non-zero, though the size and direction of any cross-sectional variation are both ambiguous *ex ante* (see Kandel and Stambaugh 1995). In the calibrations, these abnormal returns are found to exhibit cross-section patterns consistent with those for abnormal returns calculated relative to a market proxy (the value-weighted NYSE): that is, higher for small than large firms, and higher for high than low book-to-market firms. The cross-sectional dispersion in abnormal return obtained using the investor's optimal early-life portfolio, as a fraction of the dispersion obtained using the market proxy, is about 15% for both the size portfolios and for the book-to market portfolios. Increasing the investor's risk aversion from 4 to 10 makes the dispersion in abnormal return (calculated relative to the investor's optimal early-life portfolio) even bigger. Thus, hedging demand may be a partial explanation for the high expected returns on small and high book-to-market stocks. However, even with this hedging demand, the investor wants to short-sell the low book-to market portfolio to hold the high book-to-market portfolio.

The paper also provides intuition for the direction of the tilt in the risky-asset portfolio induced by the hedging demand. This intuition builds on Merton's (1973) characterization of the portfolio allocation by a multiperiod investor in a continuous time setting. The idea is that the investor's risky-asset portfolio is a combination of two portfolios. One is the tangency portfolio in mean-standard deviation space (mean-variance optimal MVO portfolio) while the other is the portfolio that is maximally correlated with the state variable (covariance-variance optimal CVO portfolio). Stevens (1998) shows that the weights in the MVO portfolio depend, in part, on the pattern of expected excess returns across the risky assets. His argument can also be used to show

that the weights in the CVO portfolio depend, in part, on the pattern of conditional covariances with the state variable across the risky assets. Thus, the tilt in the risky-asset portfolio induced by hedging demands can be characterized once the weight of CVO in the investor's optimal portfolio has been determined.

This intuition is used to better understand the tilts in the risky-asset portfolio for the investor with access to the three size portfolios or the three book-to-market portfolios. For example, the young investor holds a portfolio exhibiting a more negative covariance with dividend yield than the combination of MVO and the riskless asset with the same conditional expected return. To obtain this portfolio, the investor tilts away from the high book-to-market portfolio, even though this portfolio has the most negative covariance with dividend yield of the three portfolios. This occurs in part because the pattern of increasingly negative covariances going from the low to the high book-to-market portfolio is less pronounced than the increase in expected excess return going from the low to the high book-to-market portfolio. The high expected excess return on the high book-to-market portfolio gives it a large weight in the MVO portfolio which is cut back in the investor's risky-asset portfolio once the investor starts to care about negative covariance with dividend yield. Interestingly, the weights in the CVO portfolio are all positive using the three size or the three book-to-market portfolios. Since the young investor holds positive amounts of the CVO portfolio, this explains why the hedging demand induced by dividend yield as a predictor has the effect of making the investor's allocations less extreme.

The utility costs of using a value-weighted equity index or of ignoring predictability are also calculated. An investor using a value-weighted equity index would give up a much larger fraction of her wealth to have access to book-to-market portfolios than size portfolios. Further, an investor would give up a much larger fraction of her wealth to have access to dividend yield information than term spread information. The paper also performs a sensitivity analysis to determine which predictability parameters drive hedging demands. Both the persistence of the predictive variable and its correlation with asset returns are important for generating hedging demands. This sensitivity analysis can help explain why dividend yield as a predictive variable generates large hedging demands but term spread does not, since the former has a larger persistence parameter and return correlations of a much larger magnitude than term spread. A robustness check indicates that the

investor's allocation decision is insensitive to reasonable variation in the investor's impatience parameter. Thus, it appears that the value chosen for this parameter is not driving the results.

A number of recent papers address the issue of portfolio choice by a multi-period investor facing return predictability. Kandel and Stambaugh (1996) explore the effects of ignoring predictability in a myopic setting, while Brennan and Schwartz (1996), Brennan, Schwartz, and Lagnado (1996) and Barberis (1999) analyze numerically the impact of myopic versus dynamic decision-making. Campbell and Viceira (1998a) use log-linear approximations to solve the investor's multi-period discrete-time problem, while Kim and Omberg (1996) and Lui (1999) obtain exact analytical solutions for a range of continuous-time problems with predictability. Brandt (1998) uses the investor's Euler equations and U.S. stock returns to estimate the investor's portfolio allocation to stocks. Balduzzi and Lynch (1999) and Lynch and Balduzzi (1999) solve numerically the investor's multi-period problem with transaction costs.

However, these papers rarely allow the investor to hold multiple risky assets and none allow for the multiple risky assets to be portfolios of U.S. stocks. Brennan, Schwartz and Lagnado (1996) and Campbell and Viceira (1998b) allow investors to hold long-term bonds in addition to stocks while Ang and Bekaert (1999) consider the portfolio allocation problem when investors can invest in country funds. My paper is the first to consider portfolio choice when size-ranked and book-to-market-ranked portfolios are available to a multiperiod investor.

In another related paper, Campbell (1996) examines empirically whether cross-sectional variation in expected returns can be explained using the Euler equation from the multi-period investor's problem. Campbell assumes the existence of a representative agent and uses log-linear approximations to substitute consumption out of the Euler equation. He uses size and industry stock portfolios and bond portfolios and finds that stock market risk is the main factor determining excess returns. In part, this result follows from the high cross-sectional correlation between asset covariance with the stock market and asset covariance with news about future opportunity sets. The current paper complements Campbell's work by quantifying the direction and magnitude of the hedging demands induced by dividend yield and term spread as return predictors. It extends Campbell by considering stock portfolios formed on the basis of book-to-market.

A number of papers have made recent contributions to our understanding of the book-to-

market and size effects. Pursuing a risk-based explanation, Fama and French (1993) create mimicking portfolios formed on the basis of size and book-to-market and show that loadings with respect to these portfolios can substantially reduce the abnormal returns of extreme size and book-to-market portfolios. Ferson, Sarkissian and Simin (1999) show that this result does not always imply a risk-based explanation for the effects. Daniel and Titman (1997) present evidence that the stock characteristic (size or book-to-market) is still related to expected return after controlling for the stock's loading with respect to the mimicking portfolio. Fama and French (1995) examine whether size and book-to-market factors in fundamentals like earnings and sales can explain cross-sectional variation in expected returns. However, using portfolios formed on the basis of size and book-to-market, they find that the loadings of portfolio returns on the book-to-market factor are close to zero and do not exhibit reliable cross-sectional variation related to portfolio book-to-market. Lakonishok, Shleifer and Vishny (1994) argue that the high expected returns earned by high book-to-market (value) stocks are due to market inefficiency or suboptimal investor behavior. They show that a portfolio of value stocks is no riskier than a portfolio of "glamour" stocks along a number of dimensions. Jagannathan and Wang (1996) test a conditional CAPM with time-varying betas and a market proxy that includes the return on human capital and show that this model explains the cross-section of expected returns for size- and Beta-sorted portfolios better than the static CAPM.

Very recently, Ferson and Harvey (1999) use a set of predictive variables and reject a conditional version of the Fama-French three factor model by showing that it produces abnormal returns that vary with the predictive variables. Liew and Vassalou (1999) examine data for 10 countries and find that the book-to-market and size mimicking portfolios have incremental ability to forecast future economic growth over and above that of the market portfolio. And Lamont (1999) constructs portfolios of assets designed to track economic variables and shows that such portfolios can be useful for hedging economic risk. None of these papers examine how return predictability and a multi-period horizon affect an investor's portfolio allocation across size and book-to-market, which is the focus of the current paper.

The paper is organized as follows. Section 2 describes the investor's problem, the technique used to solve the investor's problem, and the calculation of utility cost. Section 3 calibrates asset returns to the U.S. economy. Section 4 presents the results and explanations while Section 5

concludes.

2 The Framework

2.1 The Investor's Problem

The paper considers situations where N risky assets plus a riskless asset are available for investment. The $N \times 1$ vector of risky asset returns from time t to $t+1$, \mathbf{R}_{t+1} , is either i.i.d. for all t , or predictable using a $K \times 1$ vector of instruments available at t , \mathbf{Z}_t . The risk-free rate R^f is assumed to be constant. The paper considers the optimal portfolio problem of an investor with a finite life of T periods.

The investor's preferences are assumed to be of the constant-relative-risk-aversion type (CRRA) and time separable with a rate of time preference equal to β . Expected lifetime utility is given by

$$E \left[\sum_{t=1}^T \beta^t \frac{c_t^{1-\gamma}}{1-\gamma} \mid \mathbf{Z}_1 \right], \quad (1)$$

where \mathbf{Z}_t is the vector of state variables for the investor at time t , c_t is investor's consumption at time t , and γ is the investor's relative-risk-aversion coefficient. When transaction costs are zero and returns are predictable, \mathbf{Z}_t equals the set of $K \times 1$ predictive variables. The formulation in (1) allows the investor to live until the terminal date with probability 1.

The law of motion of the investor's wealth, W , is given by

$$W_{t+1} = (W_t - c_t) \left[\mathbf{a}_t' (\mathbf{R}_{t+1} - R^f \mathbf{i}_N) + R^f \right] = W_t (1 - \kappa_t) R_{W,t+1} \quad (2)$$

where \mathbf{a}_t is the $N \times 1$ vector of portfolio weights chosen for the risky assets at t , $R_{W,t+1}$ is the portfolio return from t to $t+1$, and κ_t is the fraction of wealth consumed at t . Two problems are solved: the first allows short-selling while the second rules it out.

Given my parametric assumptions, the Bellman equation faced by the investor is given by:²

$$\frac{a(\mathbf{Z}_t, t) W_t^{1-\gamma}}{1-\gamma} = \max_{\kappa_t, \mathbf{a}_t} \left\{ \frac{\kappa_t^{1-\gamma} W_t^{1-\gamma}}{1-\gamma} + \beta (1 - \kappa_t)^{1-\gamma} W_t^{1-\gamma} \frac{1}{1-\gamma} E \left[a(\mathbf{Z}_{t+1}, t+1) R_{W,t+1}^{1-\gamma} \mid \mathbf{Z}_t \right] \right\}, \quad (3)$$

²This form of the value function derives from the CRRA utility specification in (1), and from the linearity in W of the budget constraint (2).

The optimization problem is homogeneous of degree $(1-\gamma)$ in wealth which implies that the solution is invariant to wealth. Thus the Bellman equation can be rewritten:

$$\frac{a(\mathbf{Z}_t, t)}{1-\gamma} = \max_{\kappa_t, a_t} \left\{ \frac{\kappa_t^{1-\gamma}}{1-\gamma} + \beta(1-\kappa_t)^{1-\gamma} \frac{1}{1-\gamma} E \left[a(\mathbf{Z}_{t+1}, t+1) R_{W, t+1}^{1-\gamma} | \mathbf{Z}_t \right] \right\}, \quad (4)$$

$$t = 1, \dots, T-1$$

The Bellman equation in (4) is solved by backward iteration starting with $t=T-1$ and $a(\mathbf{Z}_T, T)=1$. Thus, $a(\mathbf{Z}_t, t)$ is obtained by solving the optimization problem in (4) using $a(\mathbf{Z}_{t+1}, t+1)$ from the previous iteration.

2.2 Utility Cost Calculation.

Each set of instruments is associated with a return generating process for the return vector that reflects the predictive ability of the instruments (the C process). The marginal distribution for returns under this return generating process need not be i.i.d. However, the unconditional distribution for the return vector can be calculated. Also of interest is the return generating process (process U) that is i.i.d. with a covariance matrix equal to that for the unconditional distribution of the C process.

Two investor problems are considered for each set of instruments and associated return generating process. First, the investor can use the set of instruments \mathbf{Z}_t when making decisions at t (the conditional C problem). Alternatively, the investor can assume that the return vector follows the i.i.d. U process (the unconditional U problem).

A number of utility cost calculations are performed. When calculating the cost of using policy/process pair a rather policy/process pair b , the calculated cost represents the fraction of wealth that an investor using a would be prepared to give up to be given access to b . If one or both of the processes is a given conditional process (D, S or F), then an average cost is calculated using the unconditional distribution for the set of instruments.

3 Return Calibration.

This section describes the return and instrument data that are used, together with the quadrature approximation.

3.1 Data.

The investor's portfolio choice problem is solved for several sets of assets available to the investor:

- 1) VM set. The only available asset is the value-weighted portfolio of all assets on the NYSE, which is obtained from the CRSP Index Files.
- 2) 3M set. Three risky portfolios formed on the basis of firm size are available: M1, M2, and M3, which are value-weighted portfolios of, respectively, the smallest three, the middle four and the largest three size deciles from CRSP's Capitalization File of NYSE stocks.
- 3) 3B set. Three risky portfolios formed on the basis of firm book-to-market are available: B1, B2 and B3. These portfolios are formed from the six value-weighted portfolios SL, SM, SH, BL, BM, and BH from Fama and French (1993) and Davis, Fama and French (1999).³ The notation S (B) indicates that the firms in the portfolio are smaller (larger) than 50% of NYSE stocks. The notation L indicates that the firms in the portfolio have book-to-market ratios that place them in the bottom three deciles for all stocks; analogously, M indicates the middle four deciles and H indicates the top three deciles. The high book-to-market portfolio, B3, is an equally-weighted portfolio of SH and BH; B2 and B1 are formed similarly.
- 4) 3B2M. Six risky portfolios formed on the basis of firm book-to-market and size are available. These portfolios are the six value-weighted portfolios SL, SM, SH, BL, BM, and BH from Fama and French (1993) and Davis, Fama and French (1999).

The choice of firm characteristics used to form portfolios is predicated by the aim of achieving a wide dispersion in expected return across the portfolios. Work by Berk (1995) provides a theoretical rationale for using variables that depend on price. Both size and book-to-market satisfy this criterion. Further, empirical work by Banz (1981), Stattman (1980), and Fama and French (1992), (1993) among others finds that average return depends on both these variables even after controlling for market Beta.

The investor is allowed various sets of predictive variables when making portfolio choices:

³ I would like to thank Gene Fama and Ken French for making this data available.

- 1) U. The investor assumes returns are i.i.d.
- 2) D. The only predictive variable that the investor uses is the continuously compounded twelve-month dividend yield on the value-weighted NYSE, which is from CRSP.
- 3) S. The only predictive variable that the investor uses is the yield spread between 20-year and 1 month Treasury securities, from Ibbotson.
- 4) F. The investor uses two predictive variables, D and S.

Predictive variables are chosen to have parsimonious predictability empirically. Fama and French (1989) find that dividend yield and term spread predict distinct return components for a cross-section of asset classes. The other criterion for choosing predictive variables is some economic justification for the predictive relation. Both dividend yield and term spread move with the business cycle, making them natural predictors in a setting in which expected return moves over the business cycle.

All asset returns (including the riskfree rate) are deflated using monthly CPI inflation from CITIBASE. The data period used is from 1927:7 to 1996:11. The continuously compounded riskfree rate is estimated to be the mean of the continuously compounded one-month Treasury-bill rate from CRSP over this period, which gives a value for R^f of 0.042%.

Following Balduzzi and Lynch (1999) and Lynch and Balduzzi (1999), a VAR is estimated using OLS for each combination of assets (\mathbf{R}) and predictive variables (\mathbf{Z}). The asset return series are converted to a continuously compounded basis for the VAR; hence, \mathbf{R} is replaced by $\mathbf{r} = \ln(1+\mathbf{R})$, an $N \times 1$ vector. Without loss of generality, the predictive variables, D and S, are normalized to each be mean zero with unit variance.

The VAR is estimated assuming that \mathbf{Z}_t , a $K \times 1$ vector is the state vector at time t :

$$\mathbf{r}_{t+1} = \mathbf{a}_r + \mathbf{b}_r \mathbf{Z}_t + \mathbf{e}_{t+1} \quad (5)$$

$$\mathbf{Z}_{t+1} = \mathbf{a}_z + \mathbf{b}_z \mathbf{Z}_t + \mathbf{v}_{t+1} \quad (6)$$

where \mathbf{a}_r , $N \times 1$, and \mathbf{a}_z , $K \times 1$, are intercept vectors, \mathbf{b}_r , $N \times K$, and \mathbf{b}_z , $K \times K$, are coefficient matrices and $[\mathbf{e}_{t+1}' \ \mathbf{v}_{t+1}']'$ is an i.i.d., mean zero disturbance vector with covariance matrix $\mathbf{\Sigma}_{ev}$; the covariance matrix of \mathbf{v}_{t+1} is given by $\mathbf{\Sigma}_v$. This specification assumes that any return predictability is fully captured by \mathbf{Z}_t . The VAR implies the following expression for stock returns:

$$r_{t+1} = a_r + b_r Z_t + \boldsymbol{\eta} v_{t+1} + u_{t+1} \quad (7)$$

where $\boldsymbol{\eta}$ is a $N \times K$ vector of coefficients from a regression of e_{t+1} on v_{t+1} , and u_{t+1} is a i.i.d., mean-zero disturbance vector with covariance matrix $\boldsymbol{\Sigma}_u$ that is uncorrelated with v_{t+1} . The disturbance vector $[u_{t+1}' \ v_{t+1}']'$ is assumed to be multivariate normally distributed but with truncation for extreme realizations. Truncation is assumed so that short-selling is not ruled out by extreme realizations of e_{t+1} that have positive probability under the normal distribution but are in fact implausible.

3.2 Quadrature Approximation.

The data VAR is approximated using a variation of the Gaussian quadrature method described by Tauchen and Hussey (1991). First, Tauchen and Hussey's method is used to discretize the predictive variable vector, \mathbf{Z}_t , treating it as a first-order autoregressive process as in (6). The quadrature method is then used to calibrate a discrete distribution for the innovation \mathbf{u} . I can then calculate a discrete distribution for r_{t+1} for each $\{\mathbf{Z}_{t+1}, \mathbf{Z}_t\}$ pair from the discretization of \mathbf{Z} , since $v_{t+1} = \mathbf{Z}_{t+1} - \mathbf{a}_Z - \mathbf{b}_Z \mathbf{Z}_t$. This approach ensures that \mathbf{Z} is the only state vector. I chose a specification with 19 quadrature points for the dividend yield D , 7 for the term spread S , and 3 points for the innovations in stock returns. Balduzzi and Lynch (1999) also use this basic approach and find that the approximation is able to capture important dimensions of the predictability in the data. However, in an improvement relative to Balduzzi and Lynch, this study implements the discretization in such a way as to ensure that moments important for portfolio choice are matched exactly. In particular, the procedure matches both the conditional mean vector and the covariance matrix for log returns at all grid points of the predictive variables, as well as the unconditional volatilities of the predictive variables and the correlations of log returns with the predictive variables. Finally, the data values for $\boldsymbol{\Sigma}_{ev}$ are taken to be the covariance matrix for the associated untruncated Normal distributions when performing the quadrature approximation. But since the truncation typically uses extreme cutoffs, the resulting misstatement of $\boldsymbol{\Sigma}_{ev}$ by the approximation is likely to be small.

I find that increasing the number of grid points for stock returns from 3 to 15 has virtually no effect on the optimal portfolio weights chosen by the investor. Fifteen grid points for returns implies that the largest realization for e_{t+1} is more than six standard deviations from zero. Further,

fifteen return grid points for returns in conjunction with 19 for dividend yield as Z means that both the following events have positive probabilities for all three book-to-market portfolios: a one-month return of less than -70% and one greater than 240%. At the same time, the smallest one-month return in the data across the three book-to-market portfolios is -36% while the largest is only 61%. In addition, the investor's optimal unconditional portfolio never realizes a one-month return of less than -38%, while for the old investor's optimal conditional portfolio, the minimum one-month return is -50%, which is still far away from -100%.

There are two implications of these results. First, implausibly large deviations from the mean are needed for the possibility of negative wealth to affect the investor's portfolio choice. Second, the investor's optimal portfolio is largely unaffected by the severity of a symmetric truncation that is sufficient to ensure that the possibility of negative wealth does not drive the investor's portfolio choice. This follows from the insensitivity of the investor's portfolio choice to increases in the number of grid points, so long as portfolio value remains positive. Consequently, my results are likely to be informative of optimal portfolio choice by a CRRA investor with access to equity portfolios formed on the basis of size and book-to-market.

4 Results

This section discusses the results of the paper, starting with a comparison of moments and parameters for the data and the quadrature approximations. Then portfolio allocations are presented for an investor who lives for 20 years (240 months), has a rate of time preference β of $1/R^f$, and a relative risk aversion coefficient γ of 4. Utility cost and consumption results are also presented. Finally, several sensitivity analyses are reported and discussed.

4.1 Data and Quadrature VAR.

Table 1 reports data and quadrature parameters for the 3M asset set. Panel A reports unconditional sample moments (means, abnormal returns, standard deviations, and correlations) for the three size portfolios. Abnormal returns are estimated by regressing asset excess return on the excess return of the value-weighted NYSE (VM) taking the riskless rate R^f to be constant and equal to the sample average of 0.042%. Panel B reports data and quadrature VAR results with dividend

yield D as the state variable while Panel C reports data and quadrature VAR results with term spread S as the state variable. Table 2 reports the same data and quadrature parameters but for the 3B asset set. All results in Tables 1 and 2 are for continuously compounded returns except abnormal return which is calculated using discrete returns.

Panel A of Table 1 shows that the mean return is decreasing in firm size: ranging from 0.53% for M3 up to 0.72% for the small firm portfolio M1. Abnormal return ranges from -0.002% for M3 up to 0.129% for the small firm portfolio M1. Turning to the book-to market portfolios in Table 2, Panel A shows that the mean return is increasing in book-to-market: ranging from 0.91% for B3 down to 0.53% for the low book-to-market portfolio B1. Abnormal return is also increasing in book-to market, being 0.310% for the large book-to-market portfolio (B3) but only -0.043% for B1. These results are consistent with earlier work documenting a small firm effect and a book-to-market effect.

Panels B and C of the two tables show (consistent with earlier studies) that neither dividend yield nor term spread explain much of the empirical variation in monthly equity returns for 3M or 3B. While the R^2 s are larger for term spread, the autoregressive parameter is larger for dividend yield and the covariances between asset shocks and the predictive variable shock are much larger for dividend yield.

For the size portfolios, the covariance between asset shock and dividend yield shock is negative and its magnitude is decreasing in firm size. However, the magnitude of the correlation is increasing in firm size, reflecting the higher volatility of small firms. Turning to the book-to-market portfolios, the covariance between asset shock and dividend yield shock is also negative and its magnitude is increasing in book-to-market. However, the magnitude of the correlation is not monotonic in book-to-market.

Importantly, comparing the data and quadrature VAR parameters for either asset set, it appears that the approximation incorporates the predictability present in the data. Both the VAR regression coefficients for the returns and the covariance matrix for the return residuals are very similar for the approximation and the data, irrespective of whether D or S is being used as Z . While the conditional covariances between returns and dividend yield are higher for the approximation than the data, a close inspection reveals that this is being driven by the higher conditional volatility of D in the approximation than the data. This interpretation is confirmed by the conditional correlations

between return and D which are similar in both the approximation and the data. As discussed above, the approximation is designed to produce predictive variables with exactly unit variance and the tables confirm that this is the case. Similarly, in unreported results, the conditional means and covariances for the approximation's log asset returns exactly match those implied by the data VAR. Finally, the persistence parameter, b_z , for the approximation is typically close to but lower than that for the data. Consequently, the approximation provides conservative estimates of both the magnitude of the hedging demands and of the utility costs associated with not using the predictive variable(s).

4.2 Portfolio Allocation.

Portfolio allocations are reported for the VM, 3M and 3B asset sets in Figures 1 to 3. For each set of two graphs, the left graph shows the allocations without short-selling while the right graph shows allocations with short-selling. Each graph shows the investor's allocation as a function of the investor's age t , where $t=1$ is her first month and $t=239$ is her last. Allocations for three sets of predictive variables are plotted in each graph. The unconditional (U) portfolio allocation is plotted together with the average allocation when the investor uses dividend yield (D) or term spread (S). The averaging is performed using the unconditional distribution for the predictive variable.

4.2.1 Value-weighted Market (VM).

Figure 1 presents allocation results when the investor has access to the value-weighted index of NYSE stocks (VM). Comparing the U and D allocations, the left graph of Figure 1 shows that the average allocation to VM in the last period using D is virtually identical to the U allocation of 53%. However, as the investor gets younger, the average allocation to VM increases from 53% to 68%. This difference of 15% is the hedging demand induced by using dividend yield as a predictive variable in the absence of short-selling. The magnitude of this demand is comparable to earlier studies (see, for example, Barberis, 1999). When short selling is allowed, the direction of the hedging demand is the same but its magnitude increases to 16%: the average allocation to VM increases from 53% in the last period to 69% early in life.

In contrast, when the investor uses S, her average allocation to VM is virtually unchanged

over her life and close to the unconditional allocation, irrespective of whether short-selling is allowed. Thus, the hedging demand induced by S is small.

4.2.2 Size Portfolios (3M Set).

Figure 2 presents allocation results when the investor has access to the 3 size portfolios of NYSE stocks. The investor's allocation decision can be broken into two parts: the allocation to the risky asset portfolio (which consists of the 3 size portfolios); and the composition of the risky asset portfolio. Panel A of Figure 2 reports the average allocation to the risky asset portfolio while Panel B reports the composition of the risky asset portfolio. In particular, Panel B contains three sets of graphs, each plotting the average allocation to a size portfolio scaled by the average allocation to the risky asset portfolio.

Interestingly, Panel A of Figure 2 shows that the average allocations to the risky asset portfolio formed using M3, M2 and M1 are slightly lower than the average allocations to VM in Figure 1. For example, having access to size portfolios rather than the value weighted index when using D reduces the old investor's average allocation to stocks from 53% to 45%. Moreover, the greater flexibility afforded by the size portfolios increases slightly the investor's overall hedging demand for stocks. When the investor uses dividend yield, that demand is large, positive and of a similar magnitude to the VM case. In the absence of short-selling, the average hedging demand is 20% (compared to 15% for VM) while relaxing the short-selling constraint results in a hedging demand of 21% (compared to 16% for VM). As with VM, hedging demand is negligible when the investor uses term spread.

Turning to Panel B, the composition of the risky asset portfolio when the investor uses U is 23% in M3, 47% in M2 and 30% in the smallest stock portfolio M1: so the short-selling restriction does not bind. In the absence of return predictability, the investor places a greater fraction of her wealth in the bottom three size deciles (M1) than the top three (M3).

Allowing the investor to use the term spread variable has almost no impact on the average allocations to the size portfolios, irrespective of whether short-selling is allowed or not. However, the availability of the dividend yield variable causes the investor early in life to increase her average allocation to M3 and reduce her average allocation to M1 relative to the U case. This result is robust

to the availability or not of short-selling.

How the investor's allocations to the size portfolios change over the lifecycle depends on whether short-selling is allowed. I focus on the allocations when short-selling is allowed since the investor's first order conditions hold in all states: the pattern of the tilts are similar with or without short-selling. The right-hand side graphs indicate that the average allocation to M3 is much larger early in life than in the last period while the converse is true for M2 and M1. Thus, the hedging demand induced by dividend yield as a predictive variable causes the investor to tilt her risky-asset portfolio away from small stocks early in life.

4.2.3 Book-to-market Portfolios (3B Set).

Figure 3 presents allocation results when the investor has access to the 3 book-to-market portfolios: B3, B2 and B1. As with Figure 2, Panel A of Figure 3 reports the average allocation to the risky asset portfolio while Panel B reports the composition of the risky asset portfolio. In particular, Panel B contains three sets of two graphs, each pair plotting the average allocation to a book-to-market portfolio scaled by the average allocation to the risky asset portfolio.

Panel A of Figure 3 shows that the average allocations to the risky asset portfolio formed using B3, B2 and B1 are similar to the average allocations to VM in Figure 1 when short-selling is prohibited. In contrast, the average allocations to the risky asset are lower using the 3B asset set than VM by between 15% and 20% when short selling is allowed. However, the magnitude of the hedging demands are larger relative to those for VM, irrespective of whether short selling is allowed. In the absence of short-selling, the average hedging demand is 19% (compared to 15% for VM) while relaxing the short-selling constraint results in a hedging demand of 24% (compared to 16% for VM). As with VM and the 3M asset set, hedging demand is negligible when the investor uses term spread.

Turning to the graphs in Panel B for which short-selling is not allowed (the three left-hand side graphs), the composition of the risky asset portfolio when the investor uses U is 100% in B3, the high book-to-market portfolio. Thus, ignoring return predictability, the short-selling restriction is binding and the investor does not want to hold any of B2 or B1. When the investor is allowed to use dividend yield, the average allocation to B2 and B1 is still 0%, while the availability of term spread causes the average allocations to these portfolios to be less than 3%. Thus, the investor only

wants to hold a positive amount of the high book-to-market portfolio, and return predictability does little to alter this conclusion.

Before examining the composition of the risky-asset portfolio when short-selling is allowed, recall that the Panel B graphs scale the average holding of each risky asset by the average total investment in the risky assets. The graphs in Panel B for which short-selling is allowed (the right-hand side graphs) show that, when the investor uses U , the average allocation to B3 is more than three times the average allocation to the risky portfolio, while the average allocations to B2 and B1 are negative. Using the dividend yield to predict asset returns has little effect on average allocations to the three assets in the last period relative to the allocations using U . However, early in life, dividend yield information causes the investor to tilt away from the high book to market portfolio and towards B2 and especially B1. The average allocation to B3 as a fraction of the average allocation to the risky portfolio drops from more than 3.5 in the last period to less than 2.5 early in life. On the other hand, the average allocation to B2 goes from negative to almost zero as the investor becomes younger. Finally, the average allocation to B1 (scaled by the average allocation to the risky portfolio) goes from about -2 late in life to about -1 early in life. Thus, as the investor gets younger, hedging demand induced by dividend yield causes her to tilt her risky-asset portfolio away from B3 and toward B2 and B1. While the magnitude of this tilt is large, it is troubling that the average allocation to the small book-to-market portfolio is still so negative, even early in life.

Finally, the hedging demands induced by the term spread variable when short-selling is allowed are again negligible.

4.3 A Comparison of Optimal Portfolios to Minimum Variance Portfolios.

It is interesting to compare the investor's optimal portfolio to the minimum variance portfolio with the same expected return. This is done in the first five columns of Table 3 whose Panel A contains results for the 3M set of assets while Panel B reports for the 3B set. The first three columns report the unconditional mean, the average conditional volatility and average conditional covariance with the predictive variable of the investor's optimal portfolio. Averaging is performed using the unconditional distribution for the predictive variable. The next two columns report the average conditional volatility and average conditional covariance with the predictive variable for the

conditional minimum variance portfolio with the same conditional mean as the optimal portfolio. The first row of each panel reports portfolio moments when returns are i.i.d. (U): in this case, the conditional distribution coincides with the unconditional. The remaining 2 pairs of rows reports portfolio moments when returns are predictable using dividend yield (D) or using term spread (S), based on the investor's optimal portfolio in her first ($t=1$) and last ($t=239$) months of life.

The first row of each panel indicates that the investor's optimal portfolio is indistinguishable from minimum variance when returns are i.i.d. (U). The volatility of the optimal portfolio is equal to that of the minimum-variance portfolio with the same mean. So although returns are calibrated to be log-normal rather than normal, the investor still behaves like a mean-variance optimizer. When a predictive variable is available, the investor still behaves like a mean-variance optimizer in the last month of life ($t=239$) but now uses the conditional rather than unconditional distribution. The average conditional volatility of the optimal portfolio is identical to that of the conditionally minimum-variance portfolio with the same conditional mean. This is true for both sets of assets and either predictive variable.

It is early in life when hedging demands have the greatest effect on the investor's optimal portfolio. Turning to the first month of life ($t=1$) when the investor uses D, first note that the optimal portfolio has a higher average volatility than in the last month of life ($t=239$). In particular, when allocating across the three size portfolios (Panel A), average portfolio volatility increases from 3.50% to 4.61% as the investor gets younger. For the three book-to-market portfolios in Panel B, a similar increase is experienced (4.70% to 5.95%) as the investor goes from the last to the first month of life.

Comparing the optimal $t=1$ portfolio using D to the minimum-variance portfolio with the same conditional mean, the optimal portfolio has a higher average conditional volatility but a more negative average covariance. When allocating across the three size portfolios (Panel A), the investor is prepared to accept an average volatility of 4.61% rather than the 4.59% from the minimum-variance portfolio to have an average covariance with D of -1.08 rather than -1.03. Panel B shows that the investor makes a similar trade-off when allocating across the three book-to-market portfolios.

Thus, when multiple risky assets are available, hedging demand induced by dividend yield has two effects. First, the investor is prepared to take on more volatility given the volatility-mean

tradeoff, because this greater volatility is accompanied by a larger negative covariance with dividend yield. Second, for a given mean, the investor is prepared to accept higher volatility to obtain a larger negative covariance with dividend yield. A larger negative covariance with dividend yield is attractive because it allows the investor to hedge uncertainty about future opportunity sets. This follows because D covaries positively with future opportunity sets and the investor's risk aversion is greater than 1. These results are consistent with previous portfolio choice papers like Barberis (1999), Campbell and Viceira (1998a) and Lynch and Balduzzi (1999) with one risky asset calibrated to U.S. equity. These papers find that the negative covariance of dividend yield with U.S. equity induces a CRRA investor with risk aversion greater than 1 to hold more of the risky asset early in life than later in life.

It is interesting to examine whether this preference for negative portfolio covariance with D translates into a tilt in the investor's risky-asset portfolio toward assets with a large negative covariance with D . Tables 1 and 2 show that this covariance is most negative for the small firm portfolio (M1) in the 3M asset set, and for the high book-to-market portfolio (B3) in the 3B asset set. But despite a preference for negative portfolio covariance with D , the investor tilts her risky-asset portfolio away from M1 and B3 as she gets younger. Thus, it is not enough to look at the covariances of the assets with the state variable to determine how hedging demand affects portfolio allocations across the assets.

The reason is that for a given portfolio covariance with D , the investor wants the portfolio with the lowest conditional volatility (see Fama, 1996, for an excellent discussion of this point).⁴ The investor cares about how a change in portfolio composition affects both the portfolio's covariance with D and its volatility. In other words, the investor wants to obtain negative covariance with dividend yield at the cheapest cost in terms of increased portfolio variance. This explains why the investor may want to hold relatively more of the risky assets with the less negative covariance with dividend yield. In fact, the entire covariance matrix for the assets and the state variable, and not

⁴ Strictly speaking, the investor cares about more than just volatility since asset returns are log-normally distributed rather than normally distributed. But the investor's holdings when returns are i.i.d. or when $t=239$ indicate that, in the absence of hedging demands, the investor behaves like a mean-variance optimizer given the assets available.

just the vector of asset covariances with the state variable, affects hedging demands. This issue is explored in more detail in the next section.

Turning to the term spread variable (S), the hedging demand induced by this variable has almost no impact on the average volatility of the investor's portfolio relative to the investor's portfolio in the last period of life.

4.4 Understanding the tilt in the risky-asset portfolio weights induced by hedging demands.

4.4.1 Theory.

Treating returns and the predictive scalar Z as conditionally multivariate normal, Fama (1996) and others (discretizing Merton 1973) have shown that if Z is the relevant state variable, a young multiperiod investor cares about $E_t[R_{W,t+1}]$, $\sigma_t[R_{W,t+1}]$ and $cov_t[R_{W,t+1}, Z_{t+1}]$ where the subscript t for the moments refers to conditioning on Z_t . Moreover, the young investor's optimal portfolio is the solution, α_t^Y , to the following problem:

$$\begin{aligned} \min_{\alpha_t} \{ & \sigma_t[R_{W,t+1}] \} \quad st \\ E_t[R_{W,t+1}] - R^f &= \mu_t^Y \\ cov_t[R_{W,t+1}, Z_{t+1}] - R^f &= \delta_t^Y. \end{aligned} \quad (8)$$

where μ_t^Y is the conditional expected excess return on the young investor's optimal portfolio and δ_t^Y is the conditional covariance of the young investor's optimal portfolio with Z_{t+1} .

Let V_t be the conditional covariance matrix for the risky return vector \mathbf{R}_{t+1} , $\boldsymbol{\mu}_t$ be the conditional mean vector for the risky excess return vector $\mathbf{r}_{t+1} = \mathbf{R}_{t+1} - \mathbf{i}_N R^f$; and, $\boldsymbol{\delta}_t$ be the covariance vector between \mathbf{R}_{t+1} and Z_{t+1} . The solution to the problem in (8) is a portfolio of the following 3 portfolios:

- 1) the riskless asset: R^f .
- 2) the mean-variance optimal (MVO) portfolio:

$$\hat{\boldsymbol{\alpha}}_t^M = \lambda_t^M \mathbf{V}_t^{-1} \boldsymbol{\mu}_t, \quad (9)$$

where $\lambda_t^M = 1/(\mathbf{i}_N' \mathbf{V}_t^{-1} \boldsymbol{\mu}_t)$ and so the risky-asset weights sum to 1.

- 3) the covariance-variance optimal (CVO) portfolio:

$$\hat{\boldsymbol{\alpha}}_t^C = \lambda_t^C \mathbf{V}_t^{-1} \boldsymbol{\delta}_t, \quad (10)$$

where $\lambda_t^C = 1/(\mathbf{i}_N' \mathbf{V}_t^{-1} \boldsymbol{\delta}_t)$ and so the risky asset weights sum to 1.

The CVO portfolio is also a scalar multiple of the portfolio maximally correlated with the state variable. It follows that the risky-asset portfolio, $\hat{\boldsymbol{\alpha}}_t^Y$, held by the young investor can be written as:

$$\hat{\boldsymbol{\alpha}}_t^Y = \hat{\boldsymbol{\alpha}}_t^M (1 - \hat{a}) + \hat{\boldsymbol{\alpha}}_t^C \hat{a} \quad (11)$$

For an old investor $\hat{a} = 0$, so the tilt in the risky-asset portfolio as the investor becomes young is given by: $\hat{a} [\hat{\boldsymbol{\alpha}}_t^C - \hat{\boldsymbol{\alpha}}_t^M]$. So the tilt depends on the fraction of the young investor's risky-asset portfolio allocated to the CVO portfolio (\hat{a}) and the difference in weights between the CVO and the MVO portfolios.

One way to ascertain \hat{a} is to compare the covariance of the young investor's portfolio with Z_{t+1} to that for two portfolios, both with the same conditional expected excess return as the young investor's portfolio μ_t^Y . One is a portfolio of R^f and the MVO (denoted by MVO,Y) and the other is a portfolio of R^f and the CVO (denoted by CVO,Y). The young investor's portfolio can be written as a portfolio of MVO,Y and CVO,Y and the weight of CVO,Y in this portfolio (denoted by a) can be obtained by exploiting the fact that covariance of portfolio j with Z_{t+1} (δ_t^j) is a linear operator:

$$\delta_t^Y = \delta_t^{M,Y} (1 - a) + \delta_t^{C,Y} a \quad (12)$$

where $\delta_t^{M,Y}$ and $\delta_t^{C,Y}$ are the conditional covariances of Z_{t+1} with portfolios MVO,Y and CVO,Y respectively. Then the weight of CVO,Y in this portfolio a can be used to obtain \hat{a} by exploiting the following result proved in Appendix A:

$$\hat{a} = \frac{a (\mu_t^M / \mu_t^C)}{1 - a + a (\mu_t^M / \mu_t^C)} \quad (13)$$

where $\mu_t^M = \boldsymbol{\mu}_t' \hat{\boldsymbol{\alpha}}_t^M$ is the conditional excess return on MVO, and $\mu_t^C = \boldsymbol{\mu}_t' \hat{\boldsymbol{\alpha}}_t^C$ is the conditional excess return on CVO. Intuitively, if a portfolio of R^f and the CVO portfolio has a more negative covariance with Z_{t+1} than a portfolio of R^f and the MVO portfolio holding conditional expected return fixed, then the young investor's preference for negative covariance with Z_{t+1} will in general translate into a positive weight for CVO in the young investor's risky-asset portfolio. When the conditional

excess return on CVO exactly equals that for MVO, equation (13) confirms that this intuition holds exactly.

The second term determining the tilt is the difference in weights between the CVO and the MVO portfolios. Thus, it would be helpful to have a characterization of each of these two sets of weights. Recent work by Stevens (1998) provides the following useful characterization of the weights in the MVO portfolio:

$$\hat{\alpha}_{i,t}^M = \lambda_t^M (\mu_{i,t} - \beta_{i,t}' \boldsymbol{\mu}_{\neq i,t}) / \sigma_t[\varepsilon_i]^2 \quad (14)$$

where $\hat{\alpha}_{i,t}^M$ is the i th element of $\hat{\boldsymbol{\alpha}}_t^M$; $\boldsymbol{\mu}_{\neq i,t}$ is the (N-1)x1 subvector of $\boldsymbol{\mu}_t$ with its i th element $\mu_{i,t}$ omitted; and, $\beta_{i,t}$ is the (N-1)x1 vector of regression slope coefficients and $\sigma_t[\varepsilon_i]$ is the conditional residual volatility from a conditional regression of $R_{i,t+1}$ on $\mathbf{R}_{\neq i,t+1}$ ($\mathbf{R}_{\neq i,t+1}$ is the (N-1)x1 subvector of \mathbf{R}_{t+1} with its i th element $R_{i,t+1}$ omitted). This characterization tells us that the pattern of weights in MVO depend on the pattern of expected excess returns across the assets. Since conditional covariance of return with Z_{t+1} is a linear operator just like conditional expected excess return, the same characterization can be applied to the weights in the CVO portfolio:

$$\hat{\alpha}_{i,t}^C = \lambda_t^C (\delta_{i,t} - \beta_{i,t}' \boldsymbol{\delta}_{\neq i,t}) / \sigma_t[\varepsilon_i]^2 \quad (15)$$

where $\hat{\alpha}_{i,t}^C$ is the i th element of $\hat{\boldsymbol{\alpha}}_t^C$; $\boldsymbol{\delta}_{\neq i,t}$ is the (N-1)x1 subvector of $\boldsymbol{\delta}_t$ with its i th element $\delta_{i,t}$ omitted; and, $\beta_{i,t}$ and $\sigma_t[\varepsilon_i]$ remain defined as before. So this characterization tells us that the pattern of weights in CVO depend on the pattern of covariances with the state variable Z_{t+1} across the assets. Thus, the direction of the tilt is going to depend on the pattern of covariances with the state variable Z_{t+1} compared to the pattern of expected excess returns across the risky assets.

4.4.2 Application to the 3M Set of Assets with D as the state variable.

Table 4 reports parameters for the optimal investor portfolio at an age (t) of 1 with short-selling when the investor has access to the 3M set of assets and uses dividend yield (D) as the predictive variable (Z). Panel A reports the conditional expected excess return on the portfolio μ_t^Y together with the conditional volatility ($\sigma_t[R_{W,t+1}]$) and the conditional covariance with the predictive variable $\sigma_t[R_{W,t+1}, D_{t+1}]$ for five values of the state variable. These last two statistics are also reported for the CVO,Y and MVO,Y portfolios. Panel A also reports the conditional expected excess returns

for the CVO and MVO portfolios (μ_t^C and μ_t^M respectively), the weight of CVO, Y in the young investor's portfolio (a), and the weight of CVO in the young investor's risky-asset portfolio (\hat{a}). Consistent with the average results in Table 3, Panel A shows that the young investor's optimal portfolio has a larger negative covariance with dividend yield than the combination of R^f and MVO with the same conditional expected return. This translates into positive values for both a and \hat{a} in all 5 states. So the young investor likes negative covariance with dividend yield and incorporates that by holding a risky-asset portfolio with a positive weight in CVO.

Panel B reports the weights of the risky assets in CVO and MVO, the conditional covariances with D and the conditional expected excess returns for the risky assets, for the same five states. Since the VAR is homoscedastic, the weights of the risky assets in CVO are identical across the five states. Note that I use the conditional covariance matrix for each state in the calculations and not the covariance matrix for the residuals from the VAR reported in Table 1. Interestingly, the three CVO weights are all positive and increase monotonically going from the small firm portfolio (M1) to the large firm portfolio (M3). Comparing the risky-asset weights in CVO and MVO reveals that CVO's weight in M1 is lower than MVO's weight for all five states. Since the young investor's risky-asset portfolio has a positive weight in CVO, it follows that the investor tilts her risky-asset portfolio away from M1 as she gets younger in all five states. This is confirmed in Figure 4 which plots both the allocation to the risky-asset portfolio and the composition of the risky-asset portfolio for the same five states. The graph for M1 shows that its allocation in the risky-asset portfolio is always lower at $t=1$ than at $t=239$. Turning to M3, the CVO allocation is lower than the MVO allocation for the 4th and 7th dividend states but is higher for the 10th, 13th and 16th states. Consistent with this, the graph for M3 in Figure 4 shows that its allocation in the risky-asset portfolio is slightly lower at $t=1$ than at $t=239$ for the 4th and 7th states but is higher for the other three states.

Turning to the conditional covariances with D and the conditional expected excess returns for the risky assets, it is clear that the expected excess return pattern over the three risky assets is changing across the five states. The pattern is always monotone increasing going from M3 (biggest) to M1 (smallest), but the pattern becomes increasingly steeper going from the low D to the high D states. At the same time, the pattern of conditional covariances with D is roughly the same for all the states with the covariance becoming more negative going from M3 to M1. Interestingly, the

covariance pattern is steeper than the expected excess return pattern for the low dividend states but this changes in the high states. This observation helps explain the changes in the tilts going from the low to the high dividend yield states.

4.4.3 Application to the 3B Set of Assets with D as the state variable.

Table 5 and Figure 5 contain the same results as Table 4 and Figure 4 but for the 3B set of assets. Again, Panel A shows that the young investor holds a portfolio with a larger negative covariance than the combination of R^f and MVO with the same conditional expected return and that this translates into positive values for both a and \hat{a} . So as with the 3M set of assets, the young investor likes negative covariance with dividend yield and incorporates that by holding a risky-asset portfolio with a positive weight in CVO.

Turning to Panel B, all three CVO weights are positive, though now the B2 asset has the largest weight. Comparing the risky-asset weights in CVO and MVO reveals that CVO's weight in B3 is always lower than MVO's weight and that CVO's weight in B1 is always higher. Since the young investor's risky-asset portfolio has a positive weight in CVO, it follows that the investor tilts her risky-asset portfolio away from B3 and toward B1 as she gets younger in all five states. This is confirmed in Figure 5: the graph for B3 shows that its allocation in the risky-asset portfolio is always lower at $t=1$ than at $t=239$ while the graph for B1 shows its allocation is always higher. With respect to B2, the CVO allocation is lower than the MVO allocation for the 4th and 7th dividend states but is higher for the 10th, 13th and 16th states. Consistent with this, the graph for B2 in Figure 5 shows that its allocation in the risky-asset portfolio is slightly lower at $t=1$ than at $t=239$ for the 4th and 7th states but is higher for the other three states.

For all five states, expected excess return increases monotonically going from B1 to B3 and the conditional covariance with D becomes monotonically more negative. However, the pattern of conditional covariances is much less steep than the pattern of expected excess returns in all five states. So B3 has a more negative covariance with D than B1 but the ratio is much smaller than the ratio of the expected excess returns for B3 and B1. This observation helps explain why the young investor likes negative covariance with D but tilts away from the asset (B3) with the largest negative covariance with dividend yield. The reason is that this asset is overweighted in the old investor's

portfolio because of its high expected excess return relative to the other two assets.

4.5 Comparing the tilt in the risky-asset portfolio weights to the Merton hedging demands.

The previous section showed that the tilt in the risky-asset portfolio as the investor becomes young is given by: $\hat{a} [\hat{\alpha}_t^C - \hat{\alpha}_t^M]$. Since $\hat{\alpha}_t^M$ depends on μ_t , the tilt depends on μ_t . Using Merton (1973), the young investor's portfolio can be approximated in the following way:⁵

$$\mathbf{a}_t^Y = \hat{\alpha}_t^M \frac{1}{\gamma \lambda_t^M} + \hat{\alpha}_t^C \mathbf{v}_t \quad (16)$$

where the first term on the left hand side (denoted MVO_γ) is the single-period investor's portfolio. Hedging demand in the Merton sense is $\hat{\alpha}_t^C \mathbf{v}_t$ which does not depend on μ_t . However, in terms of understanding cross-sectional variation in expected return, it is the tilt in the risky-asset portfolio that is relevant.

The hedging demands for two sets of assets 3M and 3B are graphed in Figure 6; the left-hand graphs are for the 3M set and the right-hand are for the 3B set. The investor's average allocation to each risky asset is plotted with the averaging performed using the unconditional distribution for D . Also plotted is the single-period investor's allocation based on the first term on the left-hand side of (16) and the investor's allocation calculated from (16) with \mathbf{v}_t set equal to 1.

First note that the single-period investor's allocation from (16) is very close to the investor's actual allocation at $t=239$. The implication is that Merton's characterization of the single-period allocation in continuous time is a good approximation of the single-period allocation in the discrete time setting used here. Second, the graphs show the young investor's desire to hold more stock dominates the hedging demand. For both sets of assets, the young investor allocates more of her portfolio to all three assets on average than the old investor.

4.6 Impact of Hedging Demands on Abnormal Returns.

The previous subsections describe how the hedging demands induced by dividend yield are large and cause investors to tilt their risky-asset portfolios away from high book-to-market and small

⁵For an investor rebalancing continuously, (16) holds exactly.

stocks. Thus, hedging demand may help explain the size and book-to-market effects in the literature. However, these effects are typically measured using abnormal return which is obtained from a market model regression of excess return on the excess return of a market proxy. The small firm effect means that this abnormal return is larger for a small than large firm portfolio, while the book-to market effect means that this abnormal return is larger for a high than low book-to-market portfolio. On the other hand, hedging demands are measured above in terms of their impact on portfolio holdings, which is traditional in the portfolio choice literature.

Thus, it would be useful to assess the impact of hedging demand on abnormal return rather than on portfolio holdings. One way to do this is to run conditional regressions of asset excess return on the excess return on the investor's optimal portfolio, and examine the magnitude and cross-sectional variation in the intercepts. If the investor chooses a conditionally minimum-variance portfolio, then the intercepts will be zero; otherwise, the intercepts will be non-zero. These intercepts can be interpreted as the CAPM abnormal returns for an economy whose representative agent is the investor. The reason is as follows. If the young investor is the representative agent, then the intercepts are those that would be obtained from regressions of asset excess returns on the market excess return. However, I do not want to lean too heavily on this interpretation since the young investor's optimal risky asset holdings are often negative. Moreover, while the partial equilibrium analysis in this paper invites speculation about the role of hedging demands for the cross-section of expected returns, definitive conclusions require a general equilibrium model. Such a model is outside the scope of the current paper.

By producing an optimal portfolio that lies inside the conditional mean-variance frontier as was documented in section 4.3, the hedging demand induced by dividend yield can be expected to generate non-zero intercepts. However, as Kandel and Stambaugh (1995) demonstrate, the closeness of an inefficient portfolio to the minimum variance frontier says nothing about the relation between expected return and Beta with respect to that portfolio, unless the observations are weighted appropriately. Consequently, both the size and direction of any cross-sectional variation in these intercepts are open questions. Considerable cross-sectional variation in the intercepts would indicate that hedging demands can materially affect abnormal return calculations. Further, if the variation is in the direction implied by the size or book-to-market effects, the implication is that the hedging

demand can help explain the effect.

Abnormal returns are reported in the last three columns of Table 3 for each size portfolio in the 3M set (Panel A) and for each book-to-market portfolio in the 3B set (Panel B). The first row of each panel reports abnormal returns when returns are i.i.d. (U). The remaining 2 pairs of rows report average abnormal returns when returns are predictable using dividend yield (D) or using term spread (S). For each predictive variable, abnormal returns are reported using the investor's allocation in her first ($t=1$) and last ($t=239$) months of life.

The abnormal returns obtained using the U investor's optimal portfolio are virtually zero, for either set of assets, consistent with her optimal portfolio being minimum-variance in each case. The same is true in the last month ($t=239$) of the investor's life using D or S. While the approximation generates log-normally distributed asset returns in the limit, the investor appears to behave like a mean-variance optimizer, in the absence of any hedging demand.

However, the D row with $t=1$ in each panel shows that the hedging demand induced by D causes abnormal return to deviate from zero. Across the three size portfolios in Panel A, abnormal return ranges from -0.029% for the large firm portfolio up to -0.010% for small firm portfolio. This cross-sectional dispersion in abnormal return of 0.018% is approximately one-seventh of the dispersion of 0.131% reported in Table 1 using the VW portfolio as the market proxy. Further, abnormal return is decreasing in firm size in each case. Turning to the book-to-market portfolios in Panel B, abnormal return is higher for the high than for the low book-to-market portfolio by 0.055%, though again both are negative. Table 1 shows that abnormal return relative to VM is also increasing in book-to market, though the dispersion of 0.351% is about 7 times larger than for the optimal $t=1$ portfolio using D. So to summarize, the patterns of abnormal returns for size and book-to-market portfolios obtained relative to the investor's optimal $t=1$ portfolio are consistent with those obtained relative to the market portfolio, but not as extreme. Thus, hedging demands induced by D may be a partial explanation for the small firm and book-to-market effects.

Turning to the hedging demands induced by term spread S, the abnormal returns are close to zero for both sets of assets (3M and 3B) and exhibit almost no cross-sectional variation. Thus, the hedging demand induced by S has a negligible effect on asset abnormal returns, as well as portfolio holdings and portfolio characteristics.

Finally, related work by Roll and Ross (1994) examines how the closeness of a portfolio to the minimum-variance frontier affects the cross-sectional relation between expected asset return and asset beta with respect to the portfolio. They hold the set of assets fixed and are concerned with identifying the set of portfolios that imply a given OLS slope coefficient between expected return and beta with respect to the portfolio. However, they do not consider the magnitude or pattern of deviations from the OLS linear relation between beta and expected return. Here, the focus is on the magnitude and pattern of expected return deviations from a linear relation with beta.⁶

4.7 Volatility of portfolio allocations across states.

Table 6 reports the standard deviations of the investor's portfolio allocations when using dividend yield (D) or term spread (S). A number of results are worth noting. The weights of the risky assets in the investor's portfolio are quite volatile, with standard deviations as high as 0.85. The volatility of the total allocation to stocks is typically much lower, though still substantial. For example, when the 3B set of assets is available and D has predictive ability that is being exploited, the volatilities of the allocations to the individual stocks range from 0.85 for B2 down to 0.27 for B1, while the volatility of the total allocation to stocks is only 0.17. Since the total allocation to stocks is just the sum of the individual allocations, the low volatility for the total allocation indicates that the individual allocations must move in opposite directions going across the states.

Intuition suggests that volatility is lower when short-selling is not allowed. While this always holds for individual asset allocations, it is not always true of the total allocation when more than one risky asset is available (3M and 3M sets). Although using S as a predictive variable generates small hedging demands, Table 6 shows that the resulting portfolio allocations exhibit substantial volatility. In fact, the total allocation to stocks is almost always more volatile using S than D. Thus, S allows the investor to time the market (so the old investor's total allocation to stocks is volatile), even though the induced hedging demand is small (so the young investor's allocation is similar to the old investor's allocation). This explains why the utility cost of not using S is found to be quite high in

⁶ Note that by running market model regressions to obtain abnormal returns, I am implicitly restricting the linear relation between expected return and Beta to have a slope coefficient equal to the excess return on the portfolio rather than the OLS slope coefficient.

Table 7 (which is discussed in the next subsection). Finally, there appears to be only small volatility differences between young and old investor allocations.

4.8 Utility Cost Calculations.

Table 7 contains utility cost calculations. Panel A reports the cost of using the value-weighted index rather than one of the asset sets formed using equity characteristics. The left-hand side presents utility costs assuming no short-selling, while the right-hand side shows how the cost calculations change when short-selling is allowed. Each row refers to the set of predictive variables available to the investor, with U referring to the i.i.d. return case.⁷ Each column heading refers to the particular set of assets that the investor is comparing to VM.

A number of results in Panel A are worth mentioning. First, being able to allocate across book-to-market is much more valuable than being able to allocate across size. With dividend yield available and no short-selling, the investor is prepared to give up 22.1% of her wealth to get the 3 book-to-market portfolios but only 9.2% to get the 3 size portfolios. Second, the benefits from using book-to-market are much larger when short-selling is allowed. If using dividend yield information, access to the 3 book-to-market portfolios is worth 22.1% without short-selling, but is worth 29.3% with it. Finally, being able to use both book-to-market and size is only valuable relative to using book-to-market, if short-selling is allowed. For example, with no short-selling, the investor is prepared to give up the same fraction of her wealth (approx 22%) to get access to the three book-to-market portfolios or the six book-to-market and size based portfolios. In contrast, when short-selling is allowed, the investor is willing to give up a much larger fraction of her wealth (36.6%) to use 3B2M than to use 3B (29.3%).

Panel B reports the utility cost associated with ignoring return predictability. Each row refers to the set of predictive variables whose predictive ability is being ignored. Each column refers to the set of assets available to the investor. The left-hand side presents costs with no short-selling

⁷Row U can be obtained using the unconditional distribution for \mathbf{R} implied by the conditional distribution obtained for any of the sets of predictive variables. However, in unreported results, the utility costs are similar for U, irrespective of the set of predictive variables used, which is further evidence that the quadrature approximation is mimicking essential features of the data.

while the right-hand side reports costs when short-selling is allowed. The cost of ignoring dividend yield is typically higher than the cost of ignoring term spread. For example, if using the 3 size portfolios and short-selling is not allowed, the investor is prepared to give up 9.1% of her wealth to exploit the term spread variable, but an even larger 11.9% to exploit dividend yield information. Even so, as the example illustrates, the cost of ignoring term spread is still substantial. Thus, while term spread fails to generate large hedging demands, it allows the investor to time the market in a way that is quite valuable.

Moreover, it appears that the term spread variable has some incremental benefit over and above the dividend yield variable. For example, the investor who uses the 3 size portfolios and is not allowed to short-sell, is prepared to give up 17.9% of her wealth to simultaneously use dividend yield and term spread information. This 17.9% can be compared to the 12.2% that this investor would give up to use the dividend yield variable. This result is consistent with the view that dividend yield and term spread predict distinct components of return (see Fama and French, 1989). Finally, the cost of ignoring predictability is always higher when short-selling is allowed.

4.9 Consumption and the Set of Assets Available.

Figure 7 presents average consumption when the investor ($\gamma=4$) is allowed to short sell and has access to either the value-weighted market (VM), the three size portfolios (3M) or the 3 book-to-market portfolios (3B). The consumption number plotted is the percentage of wealth consumed multiplied by the number of periods til T plus 1. The first panel (U) shows consumption when the investor uses the unconditional distribution while the last two panels show average consumption when the investor uses dividend yield (D) or term spread (S). Each graph shows the investor's consumption as a function of the investor's age t , where $t=1$ is her first month and $t=T=240$ is her last.

Figure 7 indicates that average consumption as a fraction of wealth and after scaling by the number of months left til T plus 1 increases as the investor gets younger irrespective of the set of assets or the predictive variable being used by the investor. At T , the consumption fraction is 1 and so as the investor gets younger, the average consumption fraction multiplied by the months til $(T+1)$ becomes increasingly larger than 1. For any of the three distributional assumptions considered (i.i.d,

U; using dividend yield, D; or using term spread, S), the slope is always steepest for the 3B set of assets, and the flattest for the VM asset. Since γ is greater than 1, the income effect of an improved opportunity set would be expected to dominate the substitution effect. For this reason, the result is consistent with the finding in Table 7 discussed above that the investor with access to VM would give up a positive fraction of her wealth to have access to either the 3M asset set and an even larger positive fraction to have access to the 3B asset set. Finally, the slope for any given asset set is steepest when the investor uses dividend yield as the predictor.

4.10 Robustness Checks and Sensitivity Analysis.

4.10.1 Risk Aversion.

To assess the sensitivity of the results, the investor's risk aversion coefficient, γ , is increased to 10. The results are qualitatively similar. While the investor holds less equity, intertemporal hedging demands induced by dividend yield still cause the investor to tilt her risky-asset portfolios away from the small firm portfolio (M1) and away from the high book-to-market portfolio (B3). Moreover, abnormal returns relative to the investor's optimal early-life portfolio using D exhibit greater cross-sectional dispersion than when γ is 4. These results are consistent with intuition from Merton (1973) that the young investor, irrespective of γ , holds combinations of R^f , MVO and CVO and a higher γ reduces the investor's holding of MVO while increasing the investor's holding of CVO. Table 8 demonstrates the point. Across the three size portfolios (3M set in Panel A), the difference in abnormal return from M3 to M1 is 0.022% which is more than 15% of the dispersion in abnormal return relative to VM reported in Table 1. Similarly, across the three book-to-market portfolios (3B set in Panel B), the difference in abnormal return from B1 to B3 is 0.071% which is more than 20% of the dispersion in abnormal return relative to VM reported in Table 1. However, Table 8 indicates that hedging demands induced by the term spread (S) remain negligible when the investor's risk aversion is increased to 10. Finally, the utility cost numbers are smaller than those for the investor whose γ equals 4, but the qualitative patterns are the same.

4.10.2 Rate of Time Preference.

For all the calibrations reported above, the rate of time preference β is set equal to the

reciprocal of the riskfree rate. There is a concern that this value for β is too high since the expected return on the investor's optimal portfolio is likely to be much higher than R^f . To assess the sensitivity of the results to choice of β , the investor's problem was also solved setting the rate of time preference equal to the reciprocal of the average return on the value weighted market proxy (VM) in the data. Since the average risky-asset holding is typically less than 1, this should produce a β value that together with $1/R^f$ brackets the range of likely β values for investors. The portfolio allocation results using this low β (unreported) are virtually indistinguishable from those reported above.

4.10.3 Rebalancing Period.

Another interesting question is how varying the rebalancing period affects portfolio allocations for a twenty-year investor. One potential complication is how consumption frequency changes as the rebalancing frequency changes. To finesse this problem, I consider the portfolio allocation problem for a 20-year investor with utility over terminal wealth only. Due to computational limitations, rebalancing periods up to only three months are considered. The first allocation decision is made at time $t=1$ and the last at $t=T$ where T is 240, 120 or 80 depending on whether the investor has a one, two or three month rebalancing period. Table 9 reports portfolio allocation results as a function of rebalancing period. Results are reported for each of three sets of assets: the value-weighted market portfolio (VM); three size portfolios (3M); and, three book-to-market portfolios (3B). For each set of assets, three return distributions are considered: the unconditional (U) and the conditional using dividend yield (D) or term spread (S). Appendix B contains details of how the calibrations for the lower frequency processes are performed.

For each set of assets, Table 9 assesses the impact of rebalancing frequency at " $t=T$ " when the investor has one rebalancing period left in her life and at " $t=1$ " when the investor has 20 years left in her life. As in the earlier figures, the investor's allocation decision can be broken into two parts: the allocation to the risky-asset portfolio; and the composition of the risky-asset portfolio. "Risky" denotes the average allocation to the risky-asset portfolio while, for the asset sets with more than one asset (3M and 3B), "Composition" refers to the average allocation to each portfolio scaled by the average allocation to the risky-asset portfolio. For the conditional distributions, averaging

is performed using the unconditional distribution for the predictive variable. Short-selling is allowed.

The table shows that reducing rebalancing frequency from monthly to quarterly when returns are log-normal and i.i.d. has virtually no impact on the investor's average allocation to risky assets, irrespective of asset set. This result is consistent with Barberis (1999) who finds that increasing the buy-and-hold period for an investor facing a log-normal i.i.d. risky-asset return leaves her risky-asset allocation unaffected. Table 9 shows that the result continues to hold with multiple risky assets that are all log-normally distributed. Turning to the results when returns are predictable, reducing rebalancing frequency from monthly to quarterly causes the old investor ($t=T$) using dividend-yield to increase slightly her average allocation to risky assets. Again, Barberis finds the same increase in average risk-asset allocation when the buy-and-hold period is increased for an investor facing a single risky asset. In contrast, with term-spread predictability, the old investor's average allocation to the risky-asset portfolio is unaffected by rebalancing frequency. Interestingly, turning to the young investor, rebalancing frequency has approximately the same affect on the average risky-asset holding as that for the old investor, both in direction and magnitude. So old or young, going from monthly to quarterly rebalancing causes the average allocation to the risky-asset portfolio to increase by about 2% using D but to remain the same using S.

4.10.4 Predictability Parameters.

It would be useful to know how the predictability parameters affect the magnitude of the hedging demand. Table 10 explores the relative importance of the persistence of the predictive variable (b_z in equation (6)) and the vector of conditional correlations between log-returns and the predictive variables (denoted $\rho[r, Z]$). Each line (except the first) in the table reports allocation results for the young investor (time $t=1$) given a different set of assumptions about the generating processes (GP) for returns and the predictive variable. All generating processes for all asset sets match the unconditional log-return covariance matrix to that for the data and the slopes of the predictive regressions (b_r in equation (5)), to those for the term spread variable (S) in the data. All predictive variables are calibrated to be zero mean and unit variance. Thus, the conditional single-period return distribution is the same for all 6 generating processes, and the portfolio allocations for

the single-period investor ($t=T-1$) in the first row of the table apply to all 6 generating processes.⁸ Across generating processes I through III, return correlations with the state variable are fixed at those for term spread (S) in the data (which are small) and the persistence parameter for the predictive variable (b_Z) is allowed to vary. Across generating processes IV to VI, return correlations with the state variable are fixed at those for log dividend yield (D) in the data (which are large negative values) and again the persistence parameter for the predictive variable (b_Z) is allowed to vary. As above, the investor's allocation decision can be broken into two parts: the allocation to the risky-asset portfolio; and the composition of the risky-asset portfolio. Short-selling is allowed.

Theoretically, if the correlations between returns and the state variable are exactly zero, then the hedging demand is also zero. Generating processes I through III allow an examination of the effect of varying the persistence of the state variable when return correlations with the state variable are small but still non-zero. The average allocation to the risky-asset portfolio by the young investor is similar under generating processes I through III to that of the myopic investor reported in the first row. Even when the persistence is matched to that of dividend yield in the data (generating process III), the average allocation to the risky-asset portfolio, differs from the myopic allocation by less than 1.5% across all three asset sets. However, for the asset sets with multiple risky assets (3M and 3B), the composition to the risky-asset portfolio is affected by increasing the persistence parameter even though return correlations with the state variable are small. In contrast, when the correlations between returns and the state variable are large (processes IV-VI), varying the predictive parameter from 0 to 0.85 to 0.96 has a large impact on both the average allocation to the risky-asset portfolio and its composition. For example, when the predictive parameter increases from 0.85 to 0.96 (V vs VI), the average allocation to the risky-asset portfolio increases by at least 6% irrespective of asset set.

The correlation vector clearly has a large main effect on both the average allocation to the risky-asset portfolio and its composition. Holding the persistence parameter fixed by comparing

⁸While the return generating processes are calibrated to have the same single-period conditional distribution, it is worth checking that in fact the time T-1 allocations are the same across the discretized processes. In unreported results, the T-1 allocations for discretizations I through V are virtually identical to those in the reported in the first row of Table 10.

processes I with IV, II with V or III with VI, Table 10 shows large increases in the average allocation to the risky-asset portfolio and greatly increased tilts in composition (relative to the myopic case) when the correlation vector goes from that for S to that for D. Moreover, when the persistence parameter is larger (for example, III vs VI rather than I vs IV), switching from the S (small) to the D (large) correlation vector has a larger effect on portfolio allocation. For the 3B set, going from III to VI causes the average risky-asset allocation to increase from 34.8% to 48.1% while going from I to IV only causes an increase in this allocation of 1.7% (from 34.6% to 36.3%). The results are similar for the other two asset sets.

Thus, both larger persistence and larger correlation with returns is driving the larger hedging demands using D than using S. At the same time, these two predictability parameters interact in such a way that a larger magnitude for one means young-investor portfolio choice is more sensitive to varying the magnitude of the other.

5 Conclusion.

This paper examines portfolio allocation across stock portfolios formed on the basis of equity characteristics like size and book-to-market. The investor lives for many periods and return predictability is calibrated to U.S. data. Comparing the investor's allocation in her last period to her allocation early in life, return predictability with dividend yield causes the investor to tilt her risky-asset portfolio away from high book-to-market stocks and away from small stocks early in life. Abnormal returns relative to the investor's optimal early-life portfolio are also calculated. These abnormal returns are found to exhibit the same cross-sectional patterns as abnormal returns calculated relative to the market portfolio: higher for small than large firms, and higher for high than low book-to-market firms. Thus, hedging demand may be a partial explanation for the high expected returns documented empirically for small firms and high book-to-market firms. However, even with this hedging demand, the investor wants to short-sell the low book-to market portfolio to hold the high book-to-market portfolio.

The paper suggests that investor hedging demands may help explain cross-sectional variation in expected return. However, the analysis is partial-equilibrium and market-clearing conditions are not being imposed. In fact, the results for the book-to-market portfolios are unable to explain why

U.S. investors hold positive amounts of low book-to-market stocks. For this reason, care must be taken not to overstate the implications of the results for the cross-section of expected returns. At the same time, the hedging demands induced by dividend yield as a predictor cause the investor's risky-asset portfolios to tilt in directions that invite speculation about the ability of hedging demands to at least partially explain the book-to-market and size effects. More work of a general equilibrium nature is needed to better understand how hedging demands induced by return predictability affect equilibrium asset prices.

A number of extensions are of interest. In particular, the current study treats sample moments and regression coefficients as population moments and coefficients. It would be interesting to assess the effects of parameter uncertainty (see, for example, Kandel and Stambaugh, 1996, and Barberis, 1999). Another direction is to incorporate transaction costs into the analysis as in Balduzzi and Lynch (1999) and Lynch and Balduzzi (1999). Finally, given evidence in Jagannathan and Wang (1996) and Ferson and Harvey (1999) that conditional risk loadings are time-varying, it would be instructive to allow the conditional covariance matrix for the assets and state variables to be state-dependent.

References

- Ang, Andrew and Geert Bekaert, 1999, International asset allocation with time-varying correlations, Working Paper, Stanford University.
- Balduzzi, Pierluigi, and Anthony W. Lynch, 1999, Transaction costs and predictability: Some utility cost calculations, *Journal of Financial Economics* 52, 47-78.
- Banz, R., 1981, The relationship between return and market value of common stocks, *Journal of Financial Economics* 9, 3-18.
- Barberis, Nicholas, 1999, Investing for the long run when returns are predictable, forthcoming *Journal of Finance*.
- Berk, Jonathan B., 1995, A critique of size-related anomalies, *Review of Financial Studies* 8, 275-286.
- Brandt, Michael W., 1999, Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach, forthcoming *Journal of Finance*.
- Brennan, Michael J., and Eduardo S. Schwartz, 1996, The use of treasury bill futures in strategic asset allocation programs, Working paper (UCLA, Los Angeles, CA).
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1996, Strategic asset allocation, forthcoming the *Journal of Economic Dynamics and Control*.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373-99.
- Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298-345.
- Campbell, J. Y., and L.M. Viceira, 1998a, Consumption and Portfolio Decisions when Expected Returns are Time Varying, forthcoming *Quarterly Journal of Economics*.
- Campbell, J. Y., and L.M. Viceira, 1998b, Who should buy long-term bonds?, Working Paper, Harvard University.
- Davis, James L., Eugene F. Fama, and Kenneth R. French, 1999, Characteristics, covariances and average returns: 1929-1997, Working Paper, University of Chicago.
- Daniel, Kent, and Sheridan Titman, 1997, Evidence on the characteristics of cross-sectional variation in stock returns, *Journal of Finance* 52, 1-33.

Fama, Eugene F., 1970, Multiperiod consumption-investment decisions, *American Economic Review* 60, 163-174.

Fama, Eugene F., 1996, Multifactor portfolio efficiency and multifactor asset pricing, *Journal of Financial and Quantitative Analysis* 33, 441-465.

Fama, Eugene F., 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 F., and Kenneth R. French, 1995, Size and book-to-market factors in earnings and returns, *Journal of Finance* 50, 131-155.

Ferson, Wayne E., and Campbell R. Harvey, 1999, Conditioning Variables and the Cross-section of Stock Returns, forthcoming *Journal of Finance*.

Ferson, Wayne E., Sergei Sarkissian, and Timothy Simin, 1999, The alpha factor pricing model: A Parable, *Journal of Financial Markets* 2, 49-68.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The Conditional CAPM and the Cross-section of Expected Returns, *Journal of Finance* 51, 3-53.

Kandel, Shmuel, and Robert Stambaugh, 1995, Portfolio inefficiency and the cross-section of expected returns, *Journal of Finance* 50, 157-184.

Kandel, Shmuel, and Robert 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.

Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny, 1994, Contrarian investment, extrapolation and risk, *Journal of Finance* 49, 1541-1578.

Lamont, Owen, 1999, Economic tracking portfolios, Working Paper, University of Chicago.

Liew, Jimmy, and Maria Vassalou, 1999, Can book-to-market, size and momentum be risk factors that predict economic growth?, Working Paper, Columbia University.

- Lui, Jun, 1998, Portfolio Selection in Stochastic Environments, Working Paper, Stanford University.
- Lynch, Anthony W., and Pierluigi Balduzzi, 1999, Transaction costs and predictability: The Impact on Portfolio Choice, forthcoming Journal of Finance.
- Merton, Robert C., 1973, An intertemporal asset pricing model, *Econometrica* 41, 867-887.
- Roll, Richard, and Stephen A. Ross, 1994, On the cross-sectional relation between expected return and betas, *Journal of Finance* 49, 101-121.
- Stattman, Dennis, 1980, Book values and stock returns, *The Chicago MBA: A Journal of Selected Papers* 4, 25-45.
- Stevens, Guy, 1998, On the inverse of the covariance matrix in portfolio analysis, *Journal of Finance* 53, 1821-1827.
- Tauchen, George, and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica* 59, 317-96.

Appendix A. Derivation of (13).

Recall that μ_t^Y is the expected excess return on the young investor's optimal portfolio in state Z_t . It is possible to express the MVO,Y portfolio as a multiple of the MVO portfolio:

$$\alpha_t^{M,Y} = \hat{\alpha}_t^M \frac{\mu_t^Y}{\mu_t^M} \quad (17)$$

Similarly, it is possible to express the CVO,Y portfolio as a multiple of the CVO portfolio:

$$\alpha_t^{C,Y} = \hat{\alpha}_t^C \frac{\mu_t^Y}{\mu_t^C} \quad (18)$$

Thus, the young investor's optimal portfolio can be rewritten:

$$\alpha_t^Y = \hat{\alpha}_t^M \frac{\mu_t^Y}{\mu_t^M} (1 - a) + \hat{\alpha}_t^C \frac{\mu_t^Y}{\mu_t^C} a \quad (19)$$

which means that the young investor has $[(\mu_t^Y/\mu_t^M) (1 - a) + (\mu_t^Y/\mu_t^C) a]$ invested in her risky asset portfolio. Thus, can obtain the following expression for the young investor's risky-asset portfolio:

$$\alpha_t^Y = \hat{\alpha}_t^M \frac{(1 - a) \mu_t^Y/\mu_t^M}{(1 - a) \mu_t^Y/\mu_t^M + a \mu_t^Y/\mu_t^C} + \hat{\alpha}_t^C \frac{a \mu_t^Y/\mu_t^C}{(1 - a) \mu_t^Y/\mu_t^M + a \mu_t^Y/\mu_t^C}. \quad (20)$$

Using the definition of \hat{a} in (11) and some algebra gives (13).

Appendix B. Calibration Details for Lower Frequency Processes.

To generate a discrete distribution for the rebalancing-period returns and the predictive variable, it is useful to decompose the log asset returns into the component of log return that is explained by current and lagged values of the predictive variable (EC) and the residual i.i.d component of one-month log return that is uncorrelated with the predictive variable (RC). In terms of equation (7), the EC is the $\mathbf{a}_r + \mathbf{b}_r \mathbf{Z}_t + \boldsymbol{\eta} \mathbf{v}_{t+1}$ term and the RC is the \mathbf{u}_{t+1} term. Each component can be separately determined for the rebalancing-period log return and since the two components are independent by definition, they can be combined in a straightforward way to arrive at the rebalancing-period log return. The EC of the rebalancing-period log return is obtained by summing this component of the single-period discretization forward for two (or three) months. Its transition probability matrix can then be calculated from the transition probability matrix for the predictive variable. The discretized process for the RC term is chosen so that the covariance matrix for the $(\boldsymbol{\eta} \mathbf{v}_{t+1} + \mathbf{u}_{t+1})$ component of the rebalancing period log return matches the covariance matrix for the single-period $(\boldsymbol{\eta} \mathbf{v}_{t+1} + \mathbf{u}_{t+1})$ multiplied by the length of the rebalancing period. The discretization of the RC term is performed using quadrature approximation, with three grid points per asset.⁹ This approach carefully calculates the rebalancing-period EC directly from the single-period discretized distribution, but reduces the number of the return states by discretizing the rebalancing-period white noise component directly, rather than from the single-period white noise distribution.

Turning to the i.i.d. distribution case (U), the mean vector and unconditional covariance matrix for one-month log returns in the data are used to discretize the one-month log returns. For the asset sets with three risky assets (3M and 3B), three grid points per asset are employed while for the VM asset, nineteen grid points are used. This single-period distribution is then summed forward to obtain the rebalancing-period distribution, exploiting the fact that returns are assumed to be i.i.d.¹⁰

⁹Increasing the number of grid points per asset for the rebalancing-period residual has virtually no impact on the allocation results.

¹⁰Again, increasing the number of grid points has virtually no effect on the allocation results.

Table 1. Calibration of the 3M Asset Set: the 3 Size Portfolios

Table 1 reports moments and parameters for the 3 Size Portfolios estimated from the data and calculated for the quadrature approximation. Panel A reports unconditional sample moments for the data including abnormal return estimates from regressing excess asset returns on the excess return on the value-weighted portfolio of NYSE stocks. Panel B reports data and quadrature VAR results with dividend yield D as the state variable while Panel C reports data and quadrature VAR results with term spread S as the state variable. All results are for continuously compounded returns except abnormal return which is calculated using discrete returns.

Panel A: Data Unconditional

Asset	Mean	Average Abnormal Return	Standard Deviation, Correlation (below diagonal)		
			M3	M2	M1
M3	0.53	-0.002	5.38		
M2	0.67	0.068	0.94	7.00	
M1	0.72	0.129	0.85	0.96	8.54

Panel B: D as the only predictive variable.

VAR Coefficients	Data		Quad	
	<i>b</i>	R ²	<i>b</i>	R ²
M3	0.28	0.26	0.28	0.26
M2	0.49	0.49	0.49	0.49
M1	0.61	0.52	0.61	0.52
D	0.97	94.67	0.96	92.58

Covariance Matrix: Standard Deviation, Covariance (above diagonal), Correlation (below)									
Variable		Data				Quad			
		M3	M2	M1	D	M3	M2	M1	D
Uncond	M3	5.38	35.46	39.27	-0.88	5.38	35.46	39.27	-1.09
	M2	0.94	7.00	57.42	-0.97	0.94	7.00	57.42	-1.23
	M1	0.85	0.96	8.54	-1.03	0.85	0.96	8.54	-1.32
	D	-0.16	-0.14	-0.12	1.00	-0.20	-0.18	-0.15	1.00
Cond	M3	5.38	35.32	39.10	-1.15	5.38	35.32	39.10	-1.35
	M2	0.94	6.99	57.12	-1.44	0.94	6.99	57.12	-1.70
	M1	0.85	0.96	8.52	-1.63	0.85	0.96	8.52	-1.91
	D	-0.93	-0.89	-0.83	0.23	-0.92	-0.89	-0.82	0.27

Table 1. cont

Panel C: S as the only instrument.

VAR Coefficients	Data		Quad	
	<i>b</i>	R ²	<i>b</i>	R ²
M3	0.38	0.50	0.38	0.50
M2	0.58	0.69	0.58	0.69
M1	0.72	0.70	0.72	0.70
S	0.87	75.49	0.85	72.34

Covariance Matrix: Standard Deviation, Covariance (above diagonal), Correlation (below)									
Variable		Data				Quad			
		M3	M2	M1	S	M3	M2	M1	S
Uncond	M3	5.38	35.46	39.27	0.30	5.38	35.46	39.27	0.29
	M2	0.94	7.00	57.42	0.58	0.94	7.00	57.42	0.57
	M1	0.85	0.96	8.54	0.81	0.85	0.96	8.54	0.80
	S	0.06	0.08	0.09	1.00	0.05	0.08	0.09	1.00
Cond	M3	5.37	35.24	38.99	-0.03	5.37	35.24	38.99	-0.04
	M2	0.94	6.98	57.01	0.08	0.94	6.98	57.01	0.08
	M1	0.85	0.96	8.51	0.18	0.85	0.96	8.51	0.19
	S	-0.01	0.02	0.04	0.50	-0.01	0.02	0.04	0.53

Table 2. Calibration of the 3B Asset Set: the 3 Book-to-Market Portfolios

Table 2 reports moments and parameters for the 3 Book-to-Market Portfolios estimated from the data and calculated for the quadrature approximation. Panel A reports unconditional sample moments for the data including abnormal return estimates from regressing excess asset returns on the excess return on the value-weighted portfolio of NYSE stocks. Panel B reports data and quadrature VAR results with dividend yield D as the state variable while Panel C reports data and quadrature VAR results with term spread S as the state variable. All results are for continuously compounded returns except abnormal return which is calculated using discrete returns.

Panel A: Unconditional

Asset	Mean	Average Abnormal Return	Standard Deviation, Correlation (below diagonal)		
			B3	B2	B1
B3	0.91	0.310	7.41		
B2	0.70	0.123	0.97	6.47	
B1	0.53	-0.043	0.92	0.96	6.49

Panel B: D as the only instrument.

VAR Coefficients	Data		Quad	
	<i>b</i>	R ²	<i>b</i>	R ²
B3	0.54	0.52	0.54	0.52
B2	0.41	0.40	0.41	0.40
B1	0.41	0.40	0.41	0.40
D	0.97	94.67	0.96	92.58

Covariance Matrix: Standard Deviation, Covariance (above diagonal), Correlation (below)									
Variable		Data				Quad			
		B3	B2	B1	D	B3	B2	B1	D
Uncond	B3	7.41	46.67	44.16	-0.99	7.41	46.67	44.16	-1.26
	B2	0.97	6.47	40.25	-0.94	0.97	6.47	40.25	-1.18
	B1	0.92	0.96	6.49	-0.91	0.92	0.96	6.49	-1.15
	D	-0.13	-0.15	-0.14	1.00	-0.17	-0.18	-0.18	1.00
Cond	B3	7.39	46.45	43.94	-1.51	7.39	46.45	43.94	-1.77
	B2	0.97	6.46	40.08	-1.34	0.97	6.46	40.08	-1.58
	B1	0.92	0.96	6.48	-1.31	0.92	0.96	6.48	-1.54
	D	-0.88	-0.90	-0.88	0.23	-0.88	-0.90	-0.87	0.27

Table 2. cont

Panel C: S as the only instrument.

VAR Coefficients	Data		Quad	
	<i>b</i>	R ²	<i>b</i>	R ²
B3	0.51	0.47	0.51	0.47
B2	0.50	0.60	0.50	0.60
B1	0.56	0.73	0.56	0.73
S	0.87	75.49	0.85	72.34

Covariance Matrix: Standard Deviation, Covariance (above diagonal), Correlation (below)									
Variable		Data				Quad			
		B3	B2	B1	S	B3	B2	B1	S
Uncond	B3	7.41	46.67	44.16	0.57	7.41	46.67	44.16	0.57
	B2	0.97	6.47	40.25	0.50	0.97	6.47	40.25	0.49
	B1	0.92	0.96	6.49	0.47	0.92	0.96	6.49	0.46
	S	0.08	0.08	0.07	1.00	0.08	0.08	0.07	1.00
Cond	B3	7.39	46.41	43.88	0.13	7.39	46.41	43.88	0.14
	B2	0.97	6.45	39.97	0.06	0.97	6.45	39.97	0.06
	B1	0.92	0.96	6.46	-0.02	0.92	0.96	6.46	-0.02
	S	0.04	0.02	-0.01	0.50	0.04	0.02	-0.01	0.53

Table 3. Comparison of Optimal and Minimum-variance Portfolios: $\gamma = 4$.

Table 3 reports parameters for the optimal investor portfolio with short-selling when returns are i.i.d. (U) and when the investor uses dividend yield (D) or term spread (S) as a predictive variable (Z). The investor has a coefficient of relative risk aversion, γ , of 4. The optimal portfolio for an investor of age (t) 1 and 239 months is reported. The expected excess return on the portfolio is reported $E[r_{w,t+1}]$ together with the average conditional volatility ($\sigma_t[R_{w,t+1}]$) and the average conditional covariance with the predictive variable $\sigma_t[R_{w,t+1}, Z_{t+1}]$. These last two statistics are also reported for the minimum-variance portfolio with the same conditional mean as the investor's portfolio in each state. When returns are i.i.d, only portfolio volatility is reported. The average intercept from a conditional regression of each asset's excess return on the excess return of the investors portfolio is reported. When returns are i.i.d, this regression is unconditional.

Panel A: 3M Asset Set (the 3 Size Portfolios)

Z	t	Optimal Portfolio			Minimum-variance Portfolio		Average Intercept		
		$E[r_{w,t+1}]$	av. $\sigma_t[R_{w,t+1}]$	av. $\sigma_t[R_{w,t+1}, Z_{t+1}]$	av. $\sigma_t[R_{w,t+1}]$	av. $\sigma_t[R_{w,t+1}, Z_{t+1}]$	M3	M2	M1
U	all	0.39	3.15	na	3.15	na	0.001	0.000	-0.001
D	239	0.57	3.50	-0.77	3.50	-0.77	0.000	-0.000	-0.001
	1	0.72	4.61	-1.08	4.59	-1.03	-0.029	-0.021	-0.010
S	239	0.58	3.37	0.04	3.37	0.04	0.002	0.001	-0.001
	1	0.58	3.36	0.04	3.36	0.04	-0.002	0.001	0.004

Panel B: 3B Asset Set (the 3 Book-to-market Portfolios)

Z	t	Optimal Portfolio			Minimum-variance Portfolio		Average Intercept		
		$E[r_{w,t+1}]$	av. $\sigma_t[R_{w,t+1}]$	av. $\sigma_t[R_{w,t+1}, Z_{t+1}]$	av. $\sigma_t[R_{w,t+1}]$	av. $\sigma_t[R_{w,t+1}, Z_{t+1}]$	B3	B2	B1
U	all	0.78	4.49	na	4.49	na	-0.001	-0.000	-0.001
D	239	0.96	4.70	-0.85	4.69	-0.84	-0.002	-0.002	-0.003
	1	1.18	5.95	-1.24	5.84	-1.06	-0.088	-0.110	-0.143
S	239	0.99	4.96	0.18	4.96	0.18	-0.005	-0.004	-0.004
	1	0.99	4.94	0.18	4.94	0.18	-0.007	-0.006	-0.008

Table 4. Comparison of the Young Investor to the Mean-variance Optimizer: 3M Asset Set (the 3 Size Portfolios).

Table 4 reports parameters for the optimal investor portfolio at an age (t) of 1 with short-selling when the investor uses log dividend yield (D) as a predictive variable (Z). The investor has a coefficient of relative risk aversion, γ , of 4. Panel A reports the conditional expected excess return on the portfolio μ_t^Y together with the conditional volatility ($\sigma_t[R_{W,t+1}]$) and the conditional covariance with the predictive variable $\sigma_t[R_{W,t+1}, D_{t+1}]$. These last two statistics are also reported for the CVO,Y and MVO,Y portfolios which are combinations of R^f and CVO and of R^f and MVO respectively, with the same conditional mean as the investor's portfolio. Panel A also reports the conditional expected excess returns for the CVO and MVO portfolios (μ_t^C and μ_t^M respectively), and the weight of CVO,Y in the young investor's portfolio (a) and the weight of CVO in the young investor's risky-asset portfolio \hat{a} . Panel B reports the weights of the risky assets in CVO and MVO, the conditional covariances with D and the conditional expected excess returns for the risky assets. Each panel reports results for 5 states.

Panel A: Conditional Moments and Values for a and \hat{a} by State

D state	Optimal Young Portfolio			CVO,Y Portfolio		MVO,Y Portfolio		a	μ_t^C	μ_t^M	\hat{a}
	μ_t^Y	$\sigma_t[R_{W,t+1}]$	$\sigma_t[R_{W,t+1}, D_{t+1}]$	$\sigma_t[R_{W,t+1}]$	$\sigma_t[R_{W,t+1}, D_{t+1}]$	$\sigma_t[R_{W,t+1}]$	$\sigma_t[R_{W,t+1}, D_{t+1}]$				
4	0.16	2.39	-0.55	3.10	-0.81	2.32	-0.45	0.26	0.29	0.34	0.30
7	0.30	3.40	-0.88	3.44	-0.90	3.40	-0.87	0.34	0.49	0.49	0.34
10	0.57	4.57	-1.17	4.73	-1.23	4.56	-1.14	0.29	0.68	0.87	0.35
13	0.95	5.73	-1.42	6.22	-1.62	5.70	-1.36	0.24	0.86	1.40	0.34
16	1.45	6.78	-1.62	7.72	-2.01	6.75	-1.54	0.17	1.07	2.11	0.29

Panel B: Risky Asset Portfolio Weights, and Conditional Expected Excess Returns and Covariances with D by State

D state	CVO Portfolio Weight: $\hat{\alpha}_{i,t}^C$			MVO Portfolio Weight: $\hat{\alpha}_{i,t}^M$			Conditional Covariance with D: $\delta_{i,t}$			Cond. Expected Excess Return: $\mu_{i,t}$		
	M3	M2	M1	M3	M2	M1	M3	M2	M1	M3	M2	M1
4	0.85	0.11	0.04	2.58	-2.24	0.66	-1.40	-1.76	-1.98	0.29	0.27	0.28
7	0.85	0.11	0.04	1.20	-0.64	0.44	-1.40	-1.76	-1.99	0.47	0.58	0.68
10	0.85	0.11	0.04	0.24	0.48	0.29	-1.41	-1.77	-2.00	0.63	0.87	1.05
13	0.85	0.11	0.04	-0.53	1.37	0.16	-1.41	-1.78	-2.00	0.80	1.17	1.41
16	0.85	0.11	0.04	-1.19	2.14	0.05	-1.41	-1.78	-2.01	0.98	1.49	1.81

Table 5. Comparison of the Young Investor to the Mean-variance Optimizer: 3B Asset Set (the 3 Book to Market Portfolios).

Table 5 reports parameters for the optimal investor portfolio at an age (t) of 1 with short-selling when the investor uses log dividend yield (D) as a predictive variable (Z). The investor has a coefficient of relative risk aversion, γ , of 4. Panel A reports the conditional expected excess return on the portfolio μ_t^Y together with the conditional volatility ($\sigma_t[R_{W,t+1}]$) and the conditional covariance with the predictive variable $\sigma_t[R_{W,t+1}, D_{t+1}]$. These last two statistics are also reported for the CVO,Y and MVO,Y portfolios which are combinations of R^f and CVO and of R^f and MVO respectively, with the same conditional mean as the investor's portfolio. Panel A also reports the conditional expected excess returns for the CVO and MVO portfolios (μ_t^C and μ_t^M respectively), and the weight of CVO,Y in the young investor's portfolio (a) and the weight of CVO in the young investor's risky-asset portfolio \hat{a} . Panel B reports the weights of the risky assets in CVO and MVO, the conditional covariances with D and the conditional expected excess returns for the risky assets. Each panel reports results for 5 states.

Panel A: Conditional Moments and Values for a and \hat{a} by State

D state	Optimal Young Portfolio			CVO,Y Portfolio		MVO,Y Portfolio		a	μ_t^C	μ_t^M	\hat{a}
	μ_t^Y	$\sigma_t[R_{W,t+1}]$	$\sigma_t[R_{W,t+1}, D_{t+1}]$	$\sigma_t[R_{W,t+1}]$	$\sigma_t[R_{W,t+1}, D_{t+1}]$	$\sigma_t[R_{W,t+1}]$	$\sigma_t[R_{W,t+1}, D_{t+1}]$				
4	0.35	3.47	-0.65	6.64	-1.68	3.30	-0.41	0.19	0.35	2.14	0.59
7	0.64	4.81	-1.03	6.75	-1.71	4.65	-0.81	0.25	0.62	1.97	0.51
10	1.04	6.08	-1.35	7.79	-1.97	5.95	-1.15	0.24	0.88	2.27	0.45
13	1.55	7.25	-1.62	8.97	-2.27	7.15	-1.44	0.22	1.14	2.69	0.39
16	2.15	8.24	-1.82	10.10	-2.54	8.19	-1.69	0.16	1.42	3.21	0.30

Panel B: Risky Asset Portfolio Weights, and Conditional Expected Excess Returns and Covariances with D by State

D state	CVO Portfolio Weight: $\hat{\alpha}_{i,t}^C$			MVO Portfolio Weight: $\hat{\alpha}_{i,t}^M$			Conditional Covariance with D: $\delta_{i,t}$			Cond. Expected Excess Return: $\mu_{i,t}$		
	B3	B2	B1	B3	B2	B1	B3	B2	B1	B3	B2	B1
4	0.18	0.61	0.22	2.63	6.71	-8.34	-1.84	-1.64	-1.60	0.48	0.36	0.19
7	0.18	0.61	0.22	3.41	1.45	-3.86	-1.85	-1.64	-1.60	0.82	0.63	0.45
10	0.18	0.61	0.22	3.69	-0.44	-2.24	-1.85	-1.65	-1.60	1.15	0.88	0.70
13	0.18	0.61	0.22	3.83	-1.47	-1.37	-1.86	-1.65	-1.61	1.47	1.12	0.94
16	0.18	0.61	0.22	3.93	-2.15	-0.78	-1.87	-1.65	-1.61	1.82	1.39	1.21

Table 6. Volatility of Portfolio Allocations: $\gamma=4$

Table 3 reports standard deviations for the weights of the risky assets and the risky-asset portfolio in the investor's optimal portfolio. "Risky" refers to the investor's risky-asset portfolio. The investor rebalances monthly and has a risk aversion coefficient $\gamma=4$, intermediate consumption, and a 20-year horizon. Standard deviations are reported for the investor's portfolio at age 1 and age 239, both with and without short-selling. The table reports results for each of three sets of assets: the value-weighted market portfolio (VM); three size portfolios (3M); and three book-to-market portfolios (3B). For each set of assets, allocations using dividend yield (D) and term spread (S) are considered, with standard deviations being calculated using the unconditional distribution for the predictive variable.

Z	t	VM	3M				3B			
			Risky	M3	M2	M1	Risky	B3	B2	B1
No Short-selling										
D	239	0.25	0.19	0.14	0.18	0.12	0.25	0.25	0.00	0.00
	1	0.25	0.20	0.21	0.29	0.10	0.25	0.25	0.00	0.00
S	239	0.30	0.25	0.09	0.20	0.11	0.23	0.21	0.03	0.04
	1	0.30	0.25	0.10	0.21	0.11	0.23	0.21	0.03	0.05
Short-selling										
D	239	0.25	0.09	0.63	0.80	0.08	0.17	0.75	0.85	0.27
	1	0.26	0.13	0.61	0.80	0.07	0.20	0.76	0.85	0.28
S	239	0.33	0.22	0.23	0.37	0.08	0.36	0.33	0.21	0.48
	1	0.33	0.22	0.22	0.37	0.07	0.36	0.34	0.21	0.48

Table 7. Utility Cost Calculations: $\gamma=4$

Table 7 reports the fraction of wealth that an investor with $\gamma=4$ would be prepared to give up to be given access to the better information set or better set of assets. The left-hand side of each panel presents utility costs assuming no short-selling, while the right-hand side shows how the cost calculations change when short-selling is allowed. Panel A reports the cost of using the value-weighted index (VM) rather than one of the asset sets formed using equity characteristics for an investor. Each row refers to the set of predictive variables available to the investor, with U referring to the i.i.d. return case. Each column heading refers to the particular set of assets that the investor is comparing to VM. Panel B reports the utility cost associated with ignoring return predictability. Each row refers to the set of predictive variables whose predictive ability is being ignored. Each column refers to the set of assets available to the investor.

Panel A: Cost of Holding VM

Predictive Variables	No Short-selling			Short-selling		
	Asset Set					
	3M	3B	3B2M	3M	3B	3B2M
U	2.3	12.5	12.6	2.3	20.6	30.9
D	9.2	22.1	22.1	10.5	29.3	36.6
S	4.1	11.1	12.8	4.3	23.0	33.8
F=[D,S]	10.7	20.8		11.4	30.8	

Panel B: Cost of Using the Unconditional Distribution (U)

Predictive Variables	No Short-selling				Short-selling			
	Asset Set							
	VM	3M	3B	3B2M	VM	3M	3B	3B2M
D	7.0	12.2	13.1	13.3	7.1	13.6	14.4	14.6
S	6.8	9.1	6.1	7.9	7.0	9.5	9.9	11.4
F=[D,S]	11.3	17.9	16.0		13.0	20.1	21.4	

Table 8. Comparison of Optimal and Minimum-variance Portfolios: $\gamma = 10$.

Table 8 reports parameters for the optimal investor portfolio with short-selling when returns are i.i.d. (U) and when the investor uses dividend yield (D) or term spread (S) as a predictive variable (Z). The investor has a coefficient of relative risk aversion, γ , of 10. The optimal portfolio for an investor of age (t) 1 and 239 months is reported. The expected excess return on the portfolio is reported $E[r_{W,t+1}]$ together with the average conditional volatility ($\sigma_t[R_{W,t+1}]$) and the average conditional covariance with the predictive variable $\sigma_t[R_{W,t+1}, Z_{t+1}]$. These last two statistics are also reported for the minimum-variance portfolio with the same conditional mean as the investor's portfolio in each state. When returns are i.i.d, only portfolio volatility is reported. The average intercept from a conditional regression of each asset's excess return on the excess return of the investors portfolio is reported. When returns are i.i.d, this regression is unconditional.

Panel A: 3M Asset Set (the 3 Size Portfolios)

Z	t	Optimal Portfolio			Minimum-variance Portfolio		Average Intercept		
		$E[r_{W,t+1}]$	av. $\sigma_t[R_{W,t+1}]$	av. $\sigma_t[R_{W,t+1}, Z_{t+1}]$	av. $\sigma_t[R_{W,t+1}]$	av. $\sigma_t[R_{W,t+1}, Z_{t+1}]$	M3	M2	M1
U	all	0.16	1.26	na	1.26	na	0.001	0.000	-0.001
D	239	0.23	1.40	-0.31	1.40	-0.31	0.000	-0.000	-0.001
	1	0.30	1.97	-0.47	1.95	-0.44	-0.033	-0.023	-0.011
S	239	0.23	1.35	0.02	1.35	0.02	0.002	0.000	-0.001
	1	0.23	1.34	0.01	1.34	0.02	-0.003	0.001	0.005

Panel B: 3B Asset Set (the 3 Book-to-market Portfolios)

Z	t	Optimal Portfolio			Minimum-variance Portfolio		Average Intercept		
		$E[r_{W,t+1}]$	av. $\sigma_t[R_{W,t+1}]$	av. $\sigma_t[R_{W,t+1}, Z_{t+1}]$	av. $\sigma_t[R_{W,t+1}]$	av. $\sigma_t[R_{W,t+1}, Z_{t+1}]$	B3	B2	B1
U	all	0.31	1.79	na	1.79	na	-0.001	-0.000	-0.001
D	239	0.38	1.87	-0.34	1.87	-0.34	-0.002	-0.001	-0.002
	1	0.49	2.54	-0.54	2.47	-0.45	-0.095	-0.123	-0.166
S	239	0.40	1.98	0.07	1.98	0.07	-0.005	-0.003	-0.004
	1	0.39	1.97	0.07	1.97	0.07	-0.007	-0.006	-0.008

Table 9. Portfolio Allocations as a Function of Rebalancing Frequency for a Terminal-wealth 20-year Investor with $\gamma=4$.

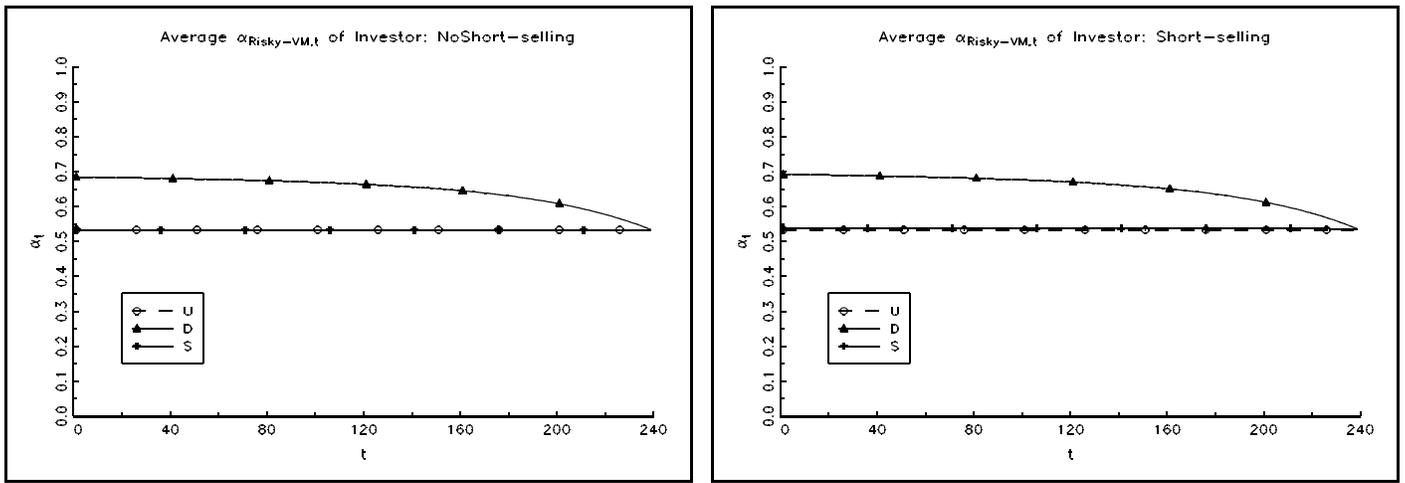
Table 9 reports allocation results when an investor with $\gamma=4$, utility over terminal wealth and a 20-year horizon rebalances at frequencies ranging from monthly to quarterly: so the rebalancing period (P) can equal 1, 2 or 3. The table reports results for each of three sets of assets: the value-weighted market portfolio (VM); three size portfolios (3M); and three book-to-market portfolios (3B). For each set of assets, the impact of rebalancing frequency is assessed at “ $t=T$ ” when the investor has one rebalancing period left in her life and at “ $t=1$ ” when the investor has 20 years left in her life. The investor’s allocation decision can be broken into two parts: the allocation to the risky-asset portfolio; and the composition of the risky-asset portfolio. “Risky” denotes the average allocation to the risky-asset portfolio while, for the asset sets with more than one asset (3M and 3B), “Composition” refers to the average allocation to each portfolio scaled by the average allocation to the risky-asset portfolio. For each set of assets, three return distributions are considered: the unconditional (U) and the conditional using dividend yield (D) or term spread (S). For the conditional distributions, averaging is performed using the unconditional distribution for the predictive variable. Short-selling is allowed.

Z	P	VM		3M							3B								
		$t=T$		$t=1$	$t=T$			$t=1$				$t=T$			$t=1$				
		Risky	Risky	Risky	Composition			Risky	Composition			Risky	Composition			Risky	Composition		
					M3	M2	M1		M3	M2	M1		B3	B2	B1		B3	B2	B1
U	1	0.534	same	0.450	0.227	0.476	0.298	same				0.345	3.712	-0.477	-2.234	same			
	2	0.534	as	0.450	0.229	0.476	0.295	as				0.346	3.690	-0.467	-2.223	as			
	3	0.534	$t=T$	0.450	0.231	0.475	0.293	$t=T$				0.347	3.669	-0.458	-2.211	$t=T$			
D	1	0.536	0.724	0.452	0.225	0.482	0.293	0.712	0.451	0.348	0.201	0.348	3.686	-0.478	-2.208	0.651	2.057	0.023	-1.080
	2	0.543	0.731	0.460	0.231	0.488	0.281	0.719	0.453	0.352	0.195	0.358	3.581	-0.451	-2.130	0.660	2.027	0.032	-1.059
	3	0.550	0.737	0.468	0.235	0.495	0.270	0.726	0.453	0.357	0.190	0.368	3.487	-0.428	-2.059	0.668	2.001	0.038	-1.040
S	1	0.536	0.536	0.452	0.225	0.480	0.295	0.459	0.282	0.452	0.266	0.346	3.692	-0.469	-2.222	0.347	3.645	-0.466	-2.179
	2	0.536	0.535	0.452	0.231	0.474	0.295	0.458	0.287	0.445	0.268	0.346	3.686	-0.465	-2.221	0.346	3.657	-0.465	-2.192
	3	0.535	0.533	0.452	0.238	0.468	0.294	0.456	0.292	0.438	0.269	0.346	3.682	-0.461	-2.221	0.344	3.667	-0.464	-2.203

Table 10. Portfolio Allocations as a Function of the Predictability Parameters for an Investor with $\gamma=4$.

Table 10 reports allocation results for an investor with $\gamma=4$, intermediate consumption, a 20-year horizon ($T=240$) and monthly rebalancing. The table reports results for each of three sets of assets: the value-weighted market portfolio (VM); three size portfolios (3M); and three book-to-market portfolios (3B). Each line (except the first) in the table reports allocation results for the young investor (time $t=1$) given a different set of assumptions about the generating processes (GP) for returns and the predictive variable. All generating processes for all asset sets match the unconditional log-return covariance matrix to that for the data and the slopes of the predictive regressions, b_r , to those for the term spread variable (S) in the data. All predictive variables are calibrated to be zero mean and unit variance. Thus, the conditional single-period return distribution is the same for all 6 generating processes, and the portfolio allocations for the single-period investor ($t=T-1$) in the first row of the table apply to all 6 generating processes. Across generating processes I to III, return correlations with the state variable ($\rho[r,Z]$) are fixed at those for term spread (S) in the data (which are small) and the persistence parameter for the predictive variable (b_z) is allowed to vary. Across generating processes IV to VI, return correlations with the state variable ($\rho[r,Z]$) are fixed at those for log dividend yield (D) in the data (which are large negative values) and again the persistence parameter for the predictive variable (b_z) is allowed to vary. The investor's allocation decision can be broken into two parts: the allocation to the risky-asset portfolio; and the composition of the risky-asset portfolio. "Risky" denotes the average allocation to the risky-asset portfolio while, for the asset sets with more than one asset (3M and 3B), "Composition" refers to the average allocation to each portfolio scaled by the average allocation to the risky-asset portfolio. Averaging is performed using the unconditional distribution for the predictive variable. Short-selling is allowed.

GP	b_r	b_z	$\rho[r,Z]$	t	VM	3M			3B				
					Risky	Risky	Composition			Risky	Composition		
							M3	M2	M1		B3	B2	B1
I-VI	S			T-1	0.536	0.452	0.225	0.480	0.295	0.346	3.692	-0.469	-2.222
I	S	0	S=small	1	0.536	0.454	0.241	0.472	0.287	0.346	3.678	-0.468	-2.210
II	S	S=0.85	S=small	1	0.536	0.459	0.280	0.453	0.267	0.347	3.647	-0.466	-2.180
III	S	D=0.96	S=small	1	0.537	0.464	0.319	0.434	0.247	0.348	3.615	-0.464	-2.151
IV	S	0	D=large neg.	1	0.566	0.488	0.272	0.452	0.276	0.363	3.522	-0.415	-2.107
V	S	S=0.85	D=large neg.	1	0.653	0.592	0.374	0.391	0.235	0.417	3.092	-0.281	-1.811
VI	S	D=0.96	D=large neg.	1	0.762	0.722	0.458	0.342	0.200	0.481	2.706	-0.166	-1.539

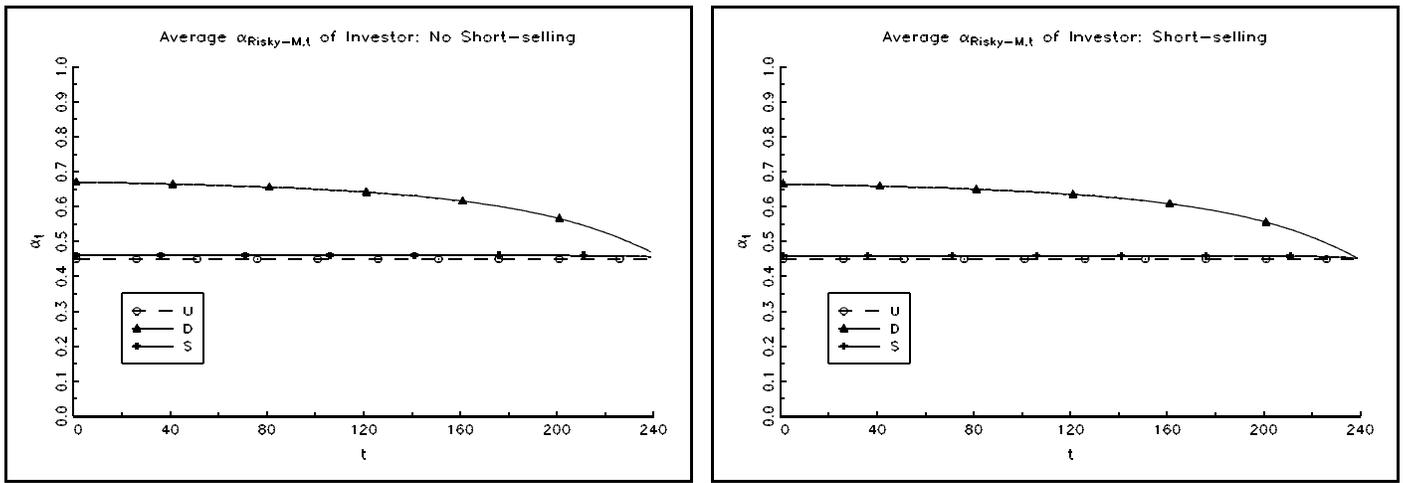


No Short-selling

Short-selling

Figure 1

Figure 1 presents allocation results when the investor ($\gamma=4$) has access to the value-weighted index of NYSE stocks (VM). The left graph shows the allocations without short-selling while the right graph shows allocations with short-selling. Each graph shows the investor's allocation as a function of the investor's age t , where $t=1$ is her first month and $t=239$ is her last. Allocations for three sets of predictive variables are plotted in each graph. The unconditional (U) portfolio allocation is plotted together with the average allocation when the investor uses dividend yield (D) or term spread (S). The averaging is performed using the unconditional distribution for the predictive variable.



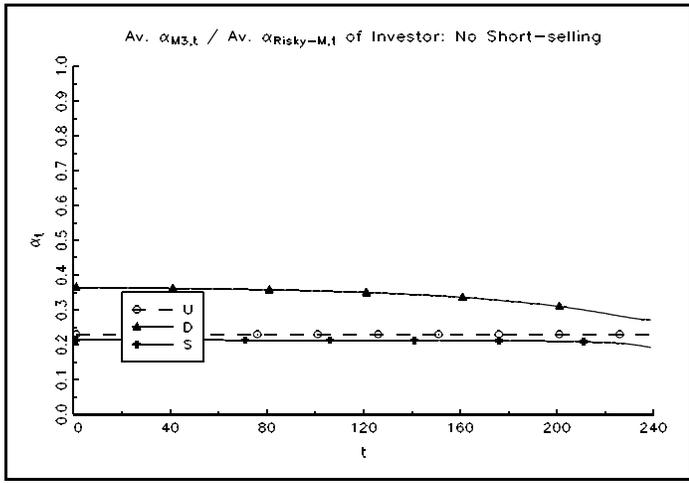
No Short-selling

Short-selling

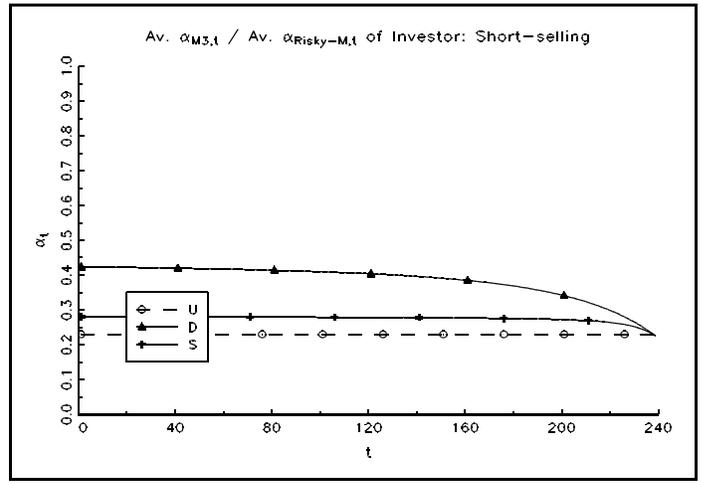
Panel A

Figure 2

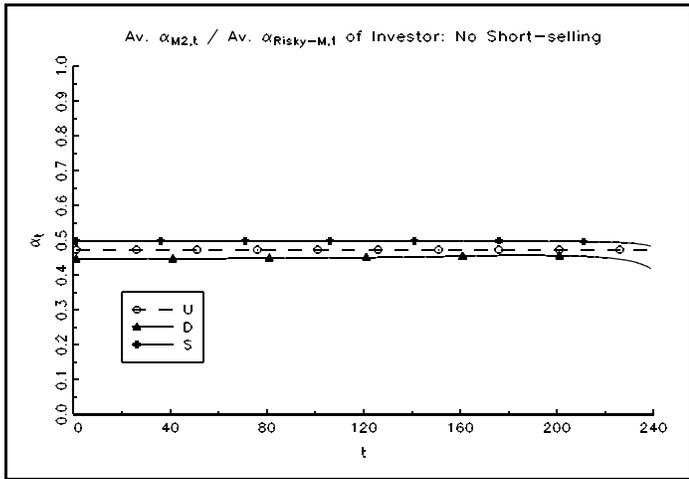
Figure 2 presents allocation results when the investor ($\gamma=4$) has access to the 3 size portfolios of NYSE stocks: M3, M2, and M1, which is the small firm portfolio. The investor's allocation decision can be broken into two parts: the allocation to the risky-asset portfolio (which consists of the 3 size portfolios); and the composition of the risky-asset portfolio. Panel A reports the average allocation to the risky-asset portfolio while Panel B reports the composition of the risky-asset portfolio. In particular, Panel B contains three sets of graphs, each plotting the average allocation to a size portfolio scaled by the average allocation to the risky-asset portfolio. For each set of two graphs, the left graph shows the allocations without short-selling while the right graph shows allocations with short-selling. Each graph shows the investor's allocation as a function of the investor's age t , where $t=1$ is her first month and $t=239$ is her last. Allocations for three sets of predictive variables are plotted in each graph. The unconditional (U) portfolio allocation is plotted together with the average allocation when the investor uses dividend yield (D) or term spread (S). The averaging is performed using the unconditional distribution for the predictive variable.



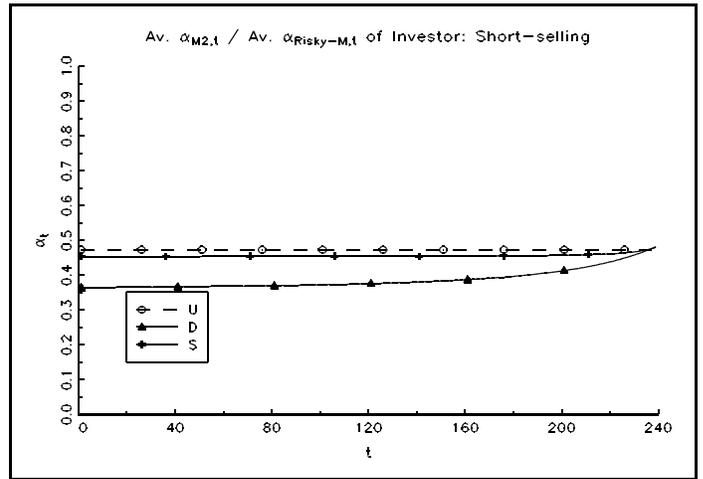
M3: No Short-selling



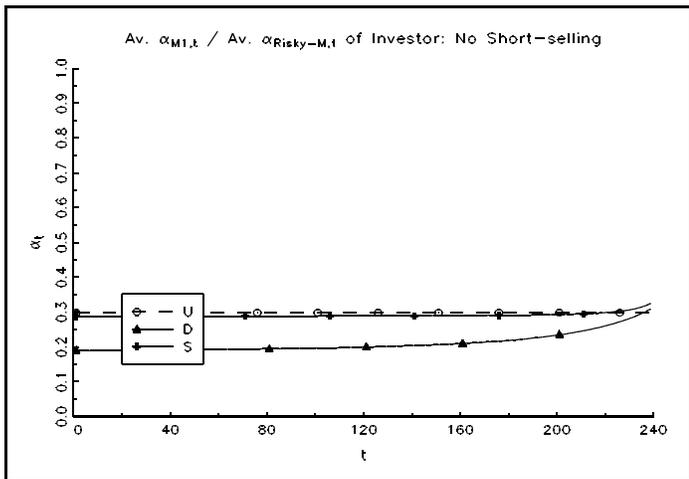
Short-selling



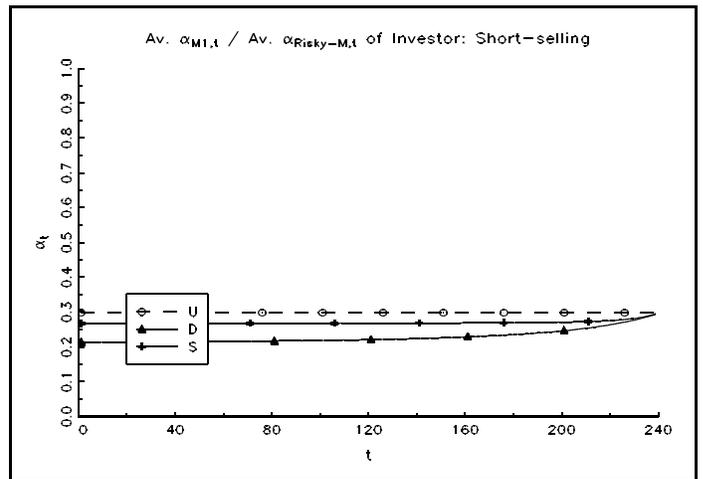
M2: No Short-selling



Short-selling



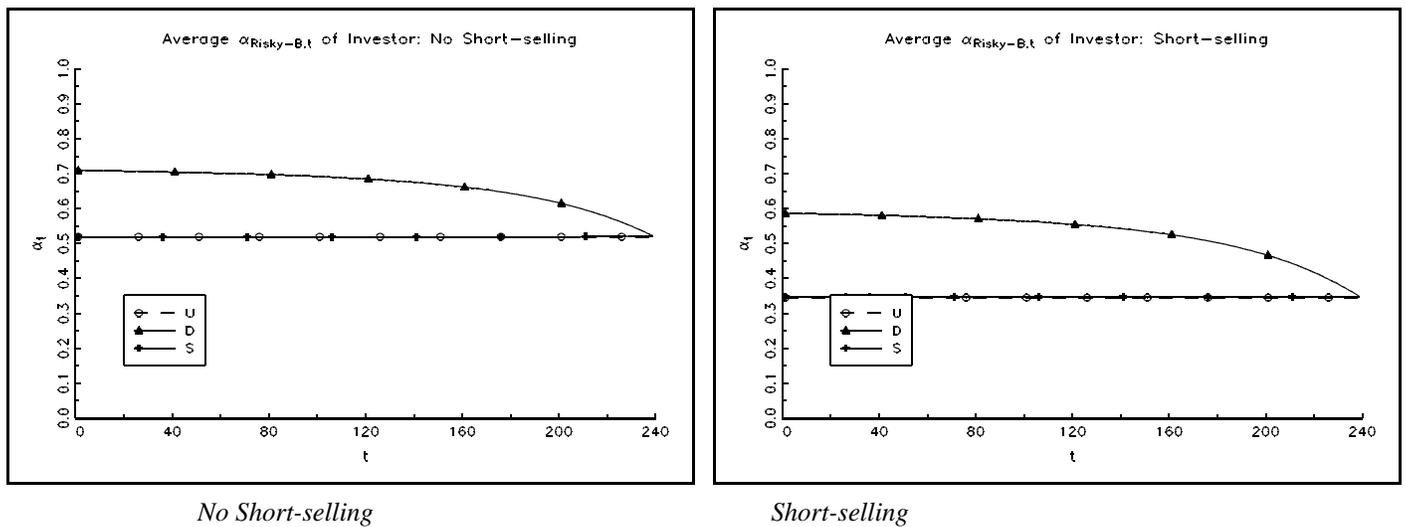
M1: No Short-selling



Short-selling

Panel B

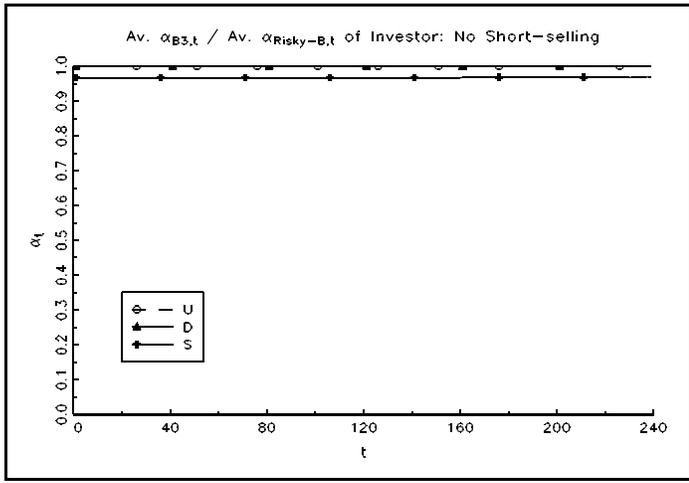
Figure 2 cont.



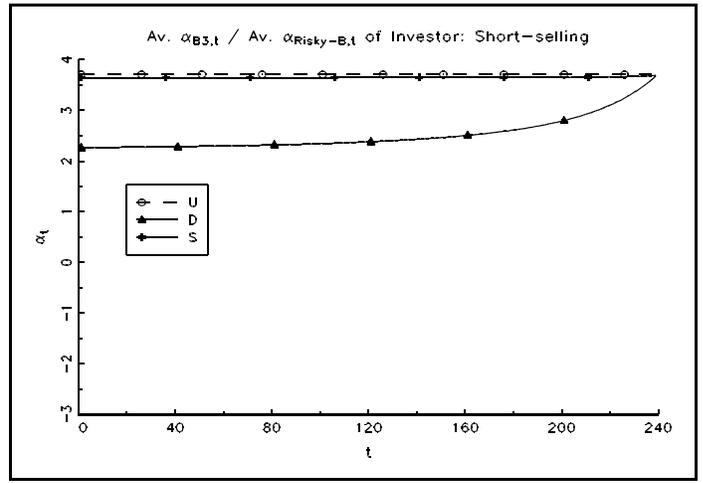
Panel A

Figure 3.

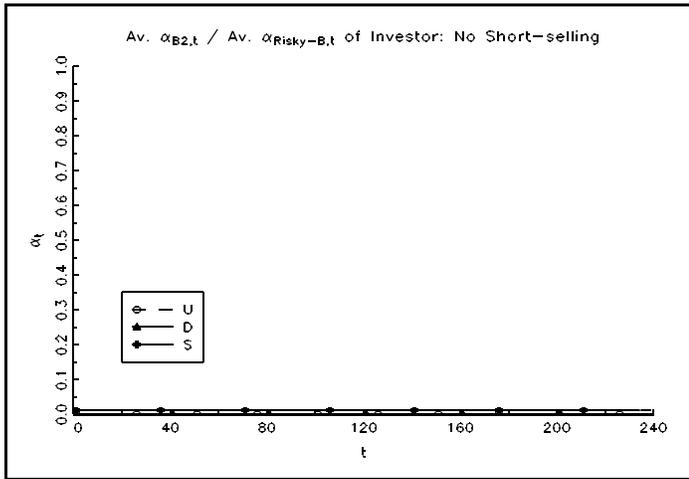
Figure 3 presents allocation results when the investor ($\gamma=4$) has access to the 3 book-to-market portfolios: B3, B2, and B1, which is the low book-to-market portfolio. The investor's allocation decision can be broken into two parts: the allocation to the risky-asset portfolio (which consists of the 3 book-to-market portfolios); and the composition of the risky-asset portfolio. Panel A reports the average allocation to the risky-asset portfolio while Panel B reports the composition of the risky-asset portfolio. In particular, Panel B contains three sets of graphs, each plotting the average allocation to a book-to-market portfolio scaled by the average allocation to the risky-asset portfolio. For each set of two graphs, the left graph shows the allocations without short-selling while the right graph shows allocations with short-selling. Each graph shows the investor's allocation as a function of the investor's age t , where $t=1$ is her first month and $t=239$ is her last. Allocations for three sets of predictive variables are plotted in each graph. The unconditional (U) portfolio allocation is plotted together with the average allocation when the investor uses dividend yield (D) or term spread (S). The averaging is performed using the unconditional distribution for the predictive variable.



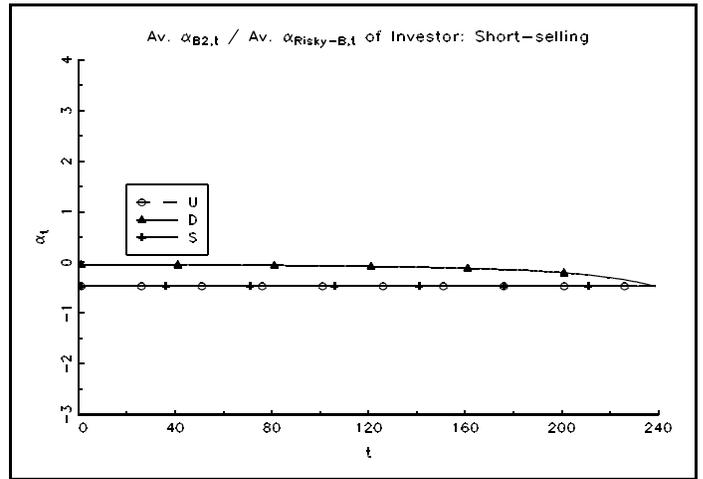
B3: No Short-selling



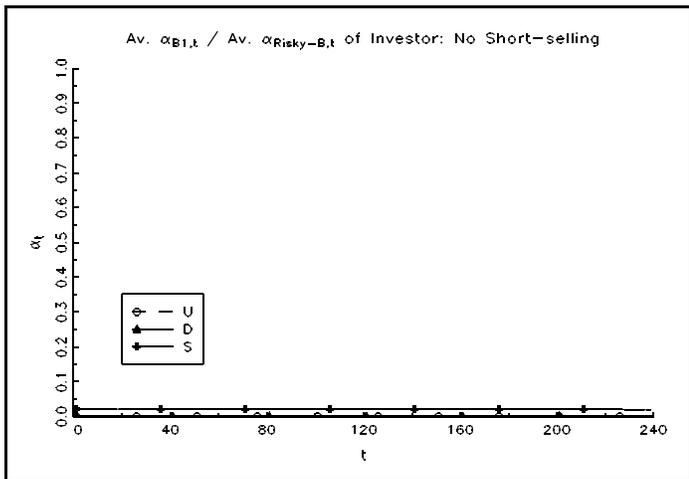
Short-selling



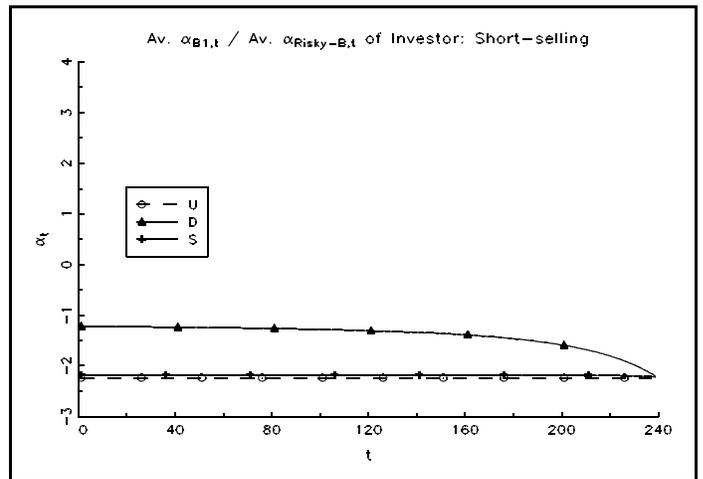
B2: No Short-selling



Short-selling



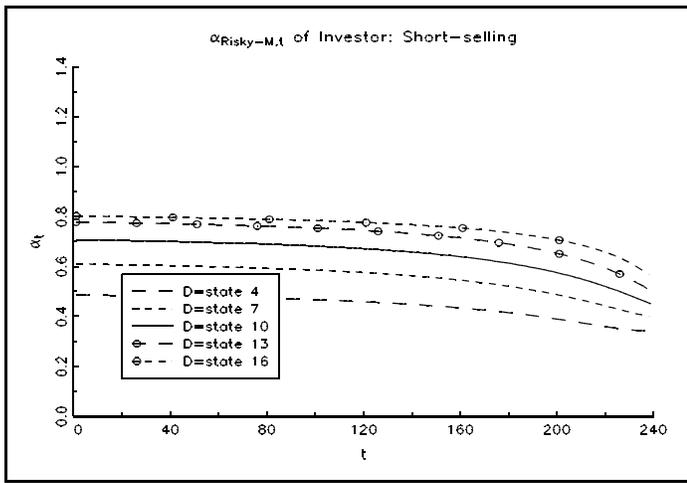
B1: No Short-selling



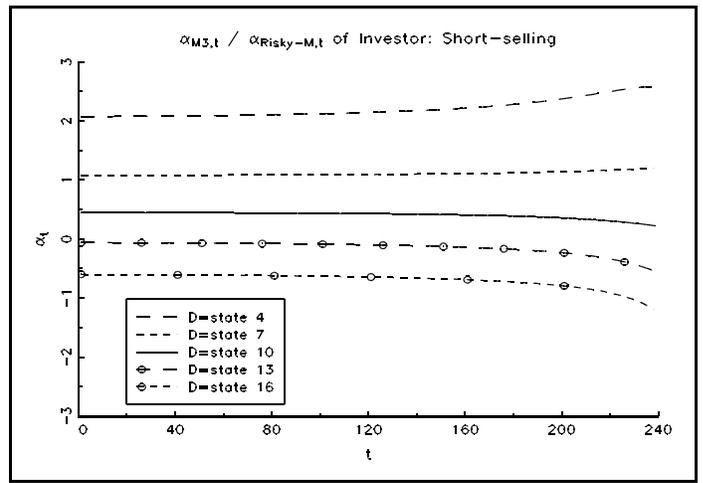
Short-selling

Panel B

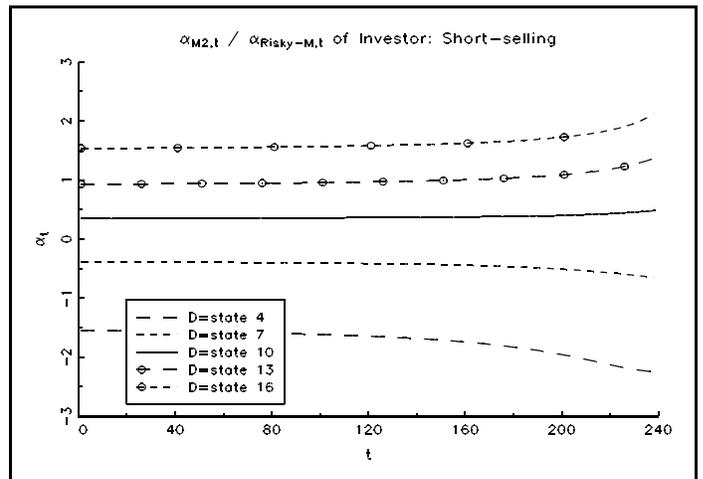
Figure 3 cont.



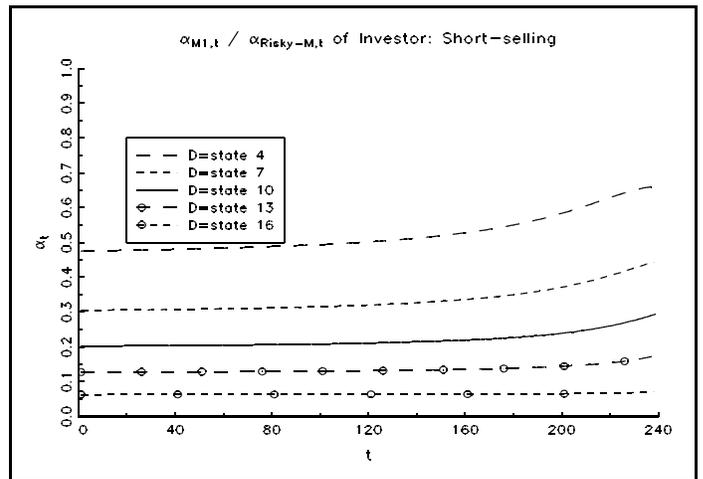
Risky-asset Portfolio



M3



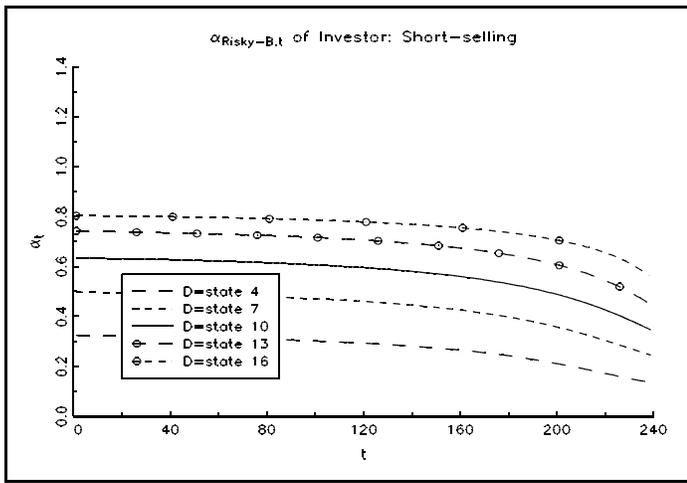
M2



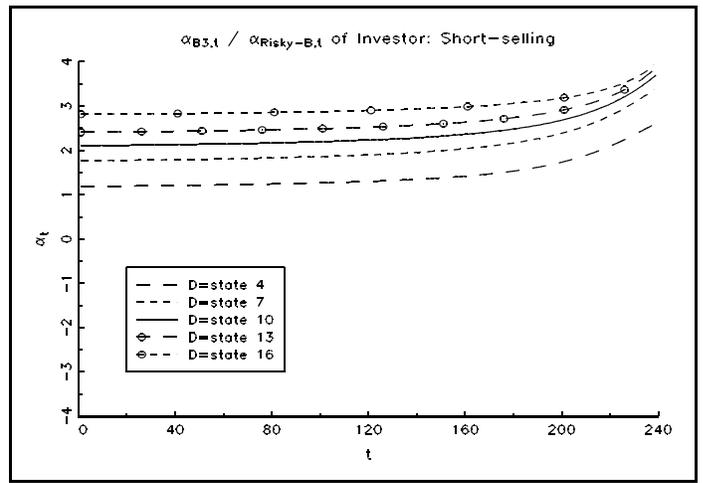
M1

Figure 4

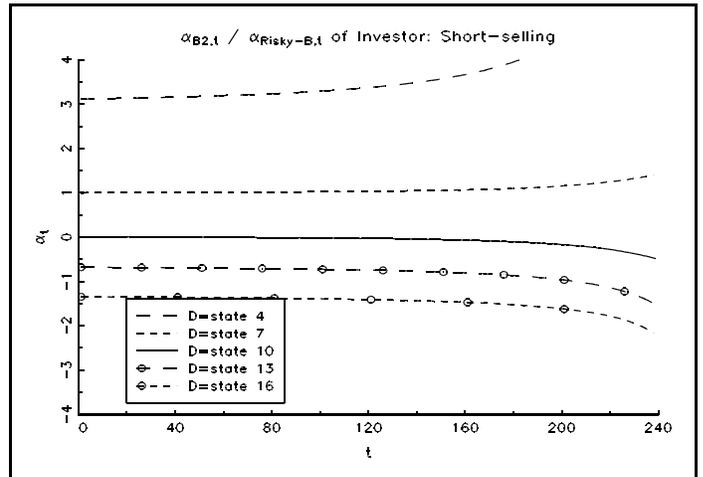
Figure 4 presents allocation results when the investor ($\gamma=4$) is allowed to short sell and has access to the 3 size portfolios : M3 (biggest), M2, and M1 (smallest). The investor uses dividend yield (D) and the plotted allocations are for five values of D. The investor's allocation decision can be broken into two parts: the allocation to the risky-asset portfolio; and the composition of the risky-asset portfolio. The left-hand graph reports the investor's allocation to the risky-asset portfolio, while the three right-hand graphs plot the weights of the three size portfolios in the risky-asset portfolio. Each graph shows the investor's allocation as a function of the investor's age t , where $t=1$ is her first month and $t=239$ is her last.



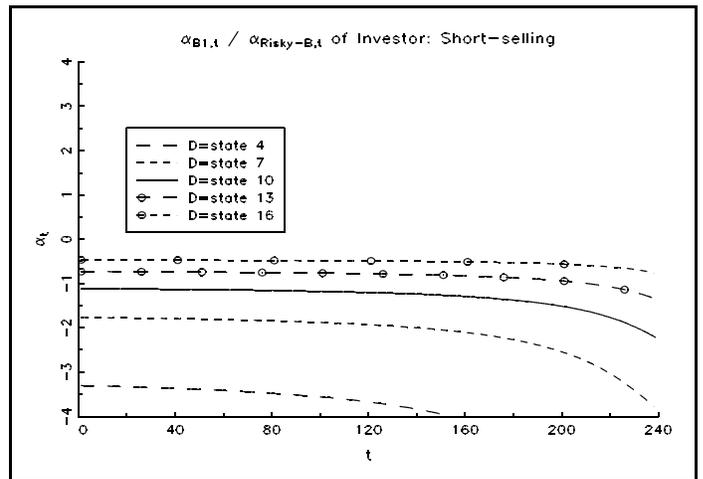
Risky-asset Portfolio



B3



B2



B1

Figure 5

Figure 5 presents allocation results when the investor ($\gamma=4$) is allowed to short sell and has access to the 3 book-to-market portfolios : B3 (highest), B2, and B1 (lowest). The investor uses dividend yield (D) and the plotted allocations are for five values of D. The investor's allocation decision can be broken into two parts: the allocation to the risky-asset portfolio; and the composition of the risky-asset portfolio. The left-hand graph reports the investor's allocation to the risky-asset portfolio, while the three right-hand graphs plot the weights of the three book-to-market portfolios in the risky-asset portfolio. Each graph shows the investor's allocation as a function of the investor's age t , where $t=1$ is her first month and $t=239$ is her last.

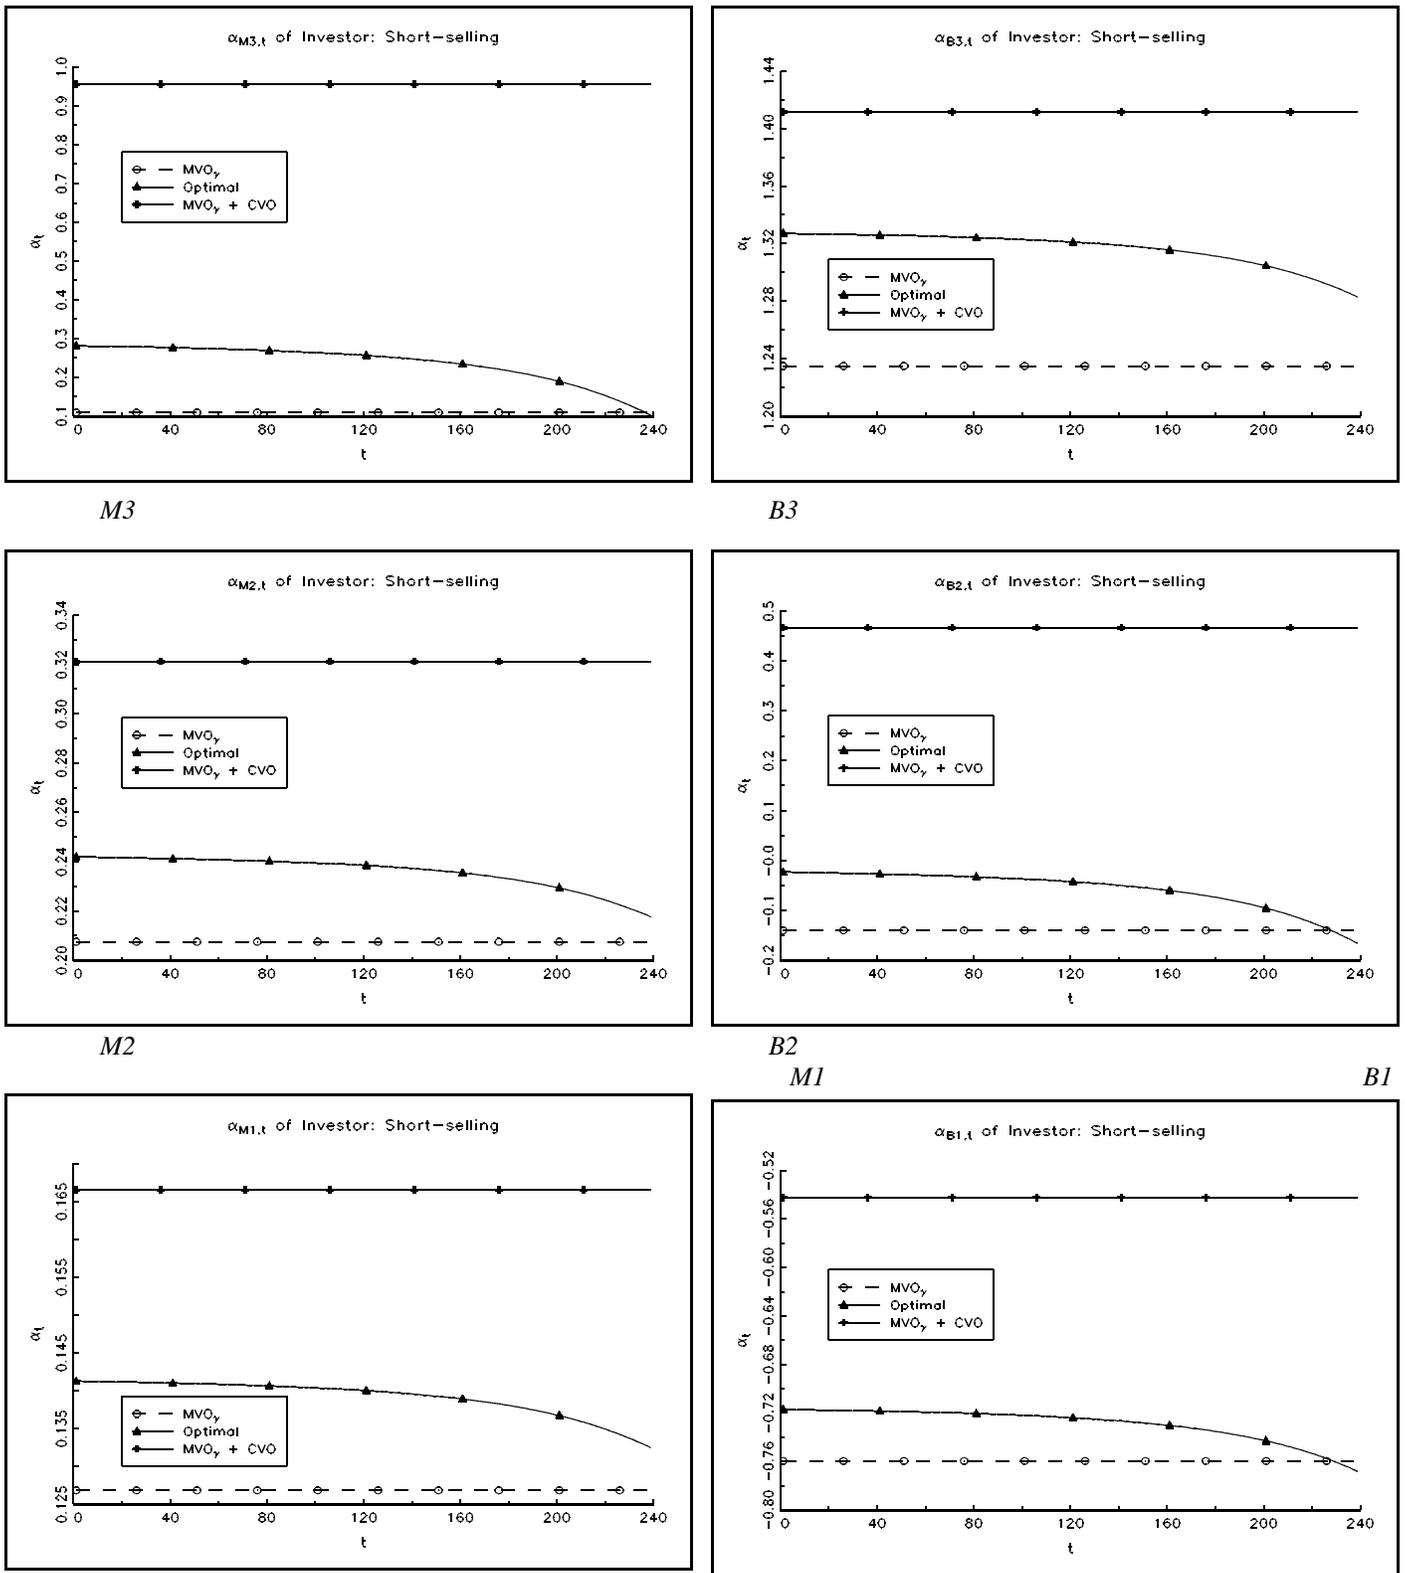
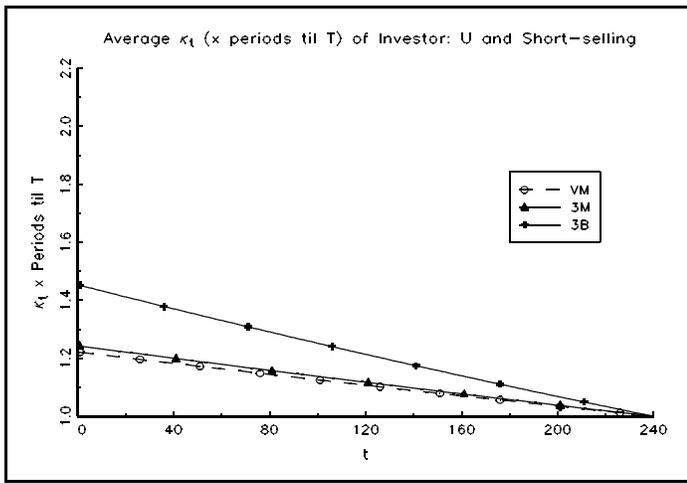
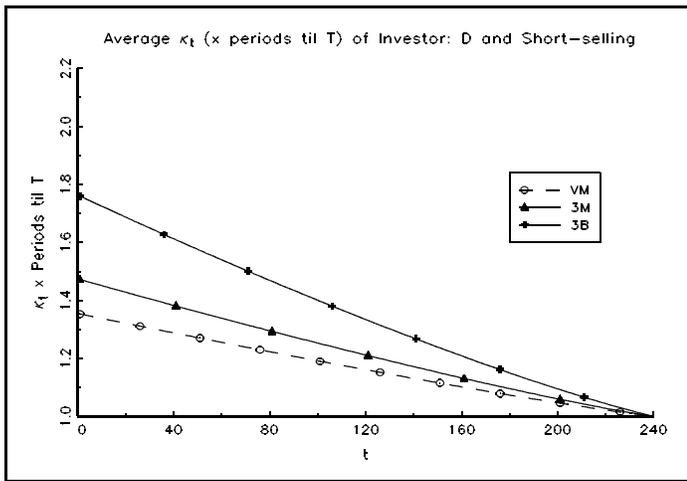


Figure 6

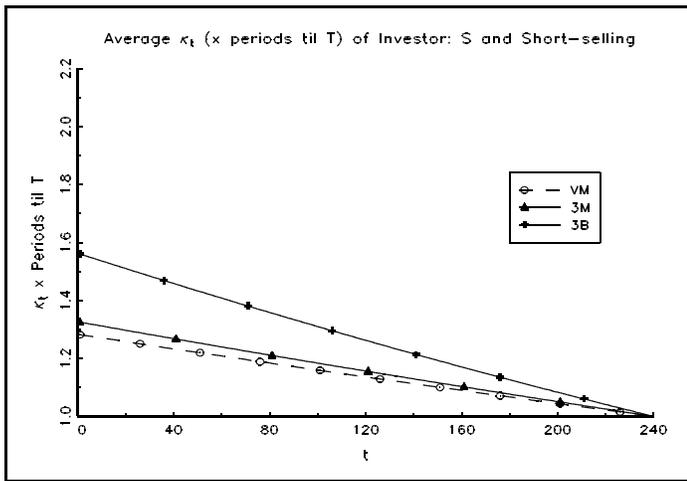
Figure 6 presents allocation results when the investor ($\gamma=4$) is allowed to short sell. The investor uses dividend yield (D). The three left-hand graphs reports the investor's allocation to the three book-to-market portfolios (B3 highest) when using the 3B asset set, while three right-hand graphs reports the investor's allocation to the three size portfolios (M1 smallest) when using the 3M asset set. Also plotted in each graph is the asset's weight in the MVO portfolio appropriately scaled by γ using Merton (1973), and the sum of this weight and the asset's weight in the CVO portfolio. Each graph shows the investor's allocation as a function of the investor's age t , where $t=1$ is her first month and $t=239$ is her last.



U



D



S

Figure 7

Figure 7 presents average consumption when the investor ($\gamma=4$) is allowed to short sell and has access to either the value-weighted market (VM), the three size portfolios (3M) or the 3 book-to-market portfolios (3B). The consumption number plotted is the percentage of wealth consumed multiplied by the number of periods til T . The first panel (U) shows consumption when the investor uses the unconditional distribution while the last two panels show average consumption when the investor uses dividend yield (D) or term spread (S). Each graph shows the investor's consumption as a function of the investor's age t , where $t=1$ is her first month and $t=T=240$ is her last.