Assessing the Economic Significance of Return Predictability: A Research Note

Walter Boudry, New York University
Philip Gray, University of Queensland

Introduction

In recent years, financial researchers have gradually accepted the notion that stock returns are partially predictable (Cochrane, 1999). Most often, the extent of return predictability is assessed from a statistical perspective, with the t-statistics and R^2's of predictive regressions guiding conclusions.[1] Statistical 'evidence' of predictability, however, does not necessarily imply economic significance.

In this paper, we assess the significance of predictor variables within an asset allocation framework. Recent research shows that the optimal allocation to risky stocks is horizon dependent if stock returns are predictable.[2] The extent of horizon effects, therefore, is a convenient metric of return predictability and our results are presented as plots of the optimal allocation to the risky asset as investment horizon increases. If a variable is useful for predicting stock returns, knowledge of that variable's value will cause a utility-maximizing investor to alter her optimal allocation. Thus, the importance of predictor variables is judged from an economic perspective, not a statistical one.

We approach the asset allocation problem using Bayesian decision theory as a unified framework for portfolio choice under uncertainty. At time T, an investor chooses weightings in a portfolio comprising one risky asset and one riskless asset, then holds this portfolio until time T+K. Portfolio weights are chosen to maximize the expected utility of terminal wealth. The key factor in the optimization problem is the distribution of forecasted cumulative returns to the risky asset over the K-period investment horizon, which is estimated using a predictive regression.

The Bayesian approach also allows the estimation risk surrounding the parameters of the predictive model to be incorporated into the density of cumulative returns. Intuitively, the more uncertain the investor is about the true degree of predictability, the less confidence she has about the model's prediction of future returns. It follows that this uncertainty ('estimation risk') is a potentially important factor in asset allocation decisions.[3]

Several prior studies have examined optimal asset allocation decisions within a Bayesian framework. Klein and Bawa (1976) compute one-period-ahead optimal allocations to one risky and one riskless asset. Since their focus is on the effect of estimation risk on portfolio decisions, they do not model return predictability, but merely assume an *i.i.d.* return-generating process. Klein and Bawa find that estimation risk causes the investor to allocate less heavily to the risky asset. However, as the sample size over which parameters are estimated increases, portfolio allocations converge to those of a classical approach that ignores estimation risk.

Kandel and Stambaugh (1996) extend this one-step-ahead problem to incorporate return predictability. Using a single predictor (dividend yield), they show that optimal allocations are highly sensitive to the most recently observed value of dividend yield even after accounting for estimation risk. Kandel and Stambaugh conclude that classical statistical tests understate the economic importance of return predictability, with p-values of 0.30 and low R^2s still containing useful information. Barberis (2000), also using dividend yield as the sole predictor, examines the optimal allocation decision of a long-horizon investor. He finds that stock returns are sufficiently predictable to induce heavier allocations to the risky asset at longer horizons. These horizon effects remain even after incorporating estimation risk.

This paper differs from prior work in several respects. First, in addition to dividend yield, our predictive regression includes two other variables which have been shown to have some predictive ability; namely, the term spread and the relative bill rate.[4] The inclusion of multiple predictors serves to highlight the potential for divergence between classical statistical significance and economic significance. It also raises the question of how to assess the significance of each predictor. Our second contribution, therefore, is to illustrate how the economic significance metric within the asset allocation framework allows a meaningful comparison of the relative importance of each predictor. We assess the relative importance of each predictor by measuring the sensitivity of optimal allocations to the last observed value of the predictor variable. If a variable can predict stock returns, the optimal asset allocation will be sensitive to the variable's most recent value. Finally, we document the economic significance of predictability for Australian equity returns. In light of the results of Barberis (2000) for US data, we compare and contrast the degree of predictability across the two markets, and document some interesting differences.

Our major findings can be summarized as follows. Ignoring estimation risk, there is sufficient return predictability to induce a risk-averse investor to allocate 9% more to risky stocks for a 10-year investment horizon compared to a 1-year horizon. These horizon effects are notably smaller than those

reported by Barberis (2000) for US data, despite the fact that he uses only a single predictor. After incorporating estimation risk, the additional dimension of uncertainty negates the horizon effects from predictability. In fact, a long-horizon investor allocates slightly less to risky stocks than his short-horizon counterpart. This finding contrasts with Barberis, who reports that estimation risk dampens, but does not eliminate, the positive horizon effects.

Of our three predictors, dividend yield is easily the most important for optimal portfolio allocations. The allocation to risky stocks is highly sensitive to the most recent value of dividend yield. While this finding accords with classical inference, the same cannot be said for the other predictors. The t-statistics imply the term spread and the relative bill rate have no predictive ability. Our results, however, show that classical statistical significance fails to capture the economic importance of these variables – over short to medium term investment horizons, optimal portfolio allocations are sensitive to the most recent value of term spread, and to a lesser extent, the relative bill rate.

The remainder of this paper is structured as follows. Section 2 outlines the asset allocation decision facing our investor and discusses how estimation risk can be either ignored or incorporated depending on the method of constructing the density of future cumulative stock returns. Section 3 presents a multivariate regression in which stock returns are predicted using three variables. The horizon effects induced by predictability, both with and without estimation risk, are documented. The relative impact of each predictor on the asset allocation decision is examined in Section 4 by measuring the sensitivity of the optimal portfolio allocation to 'shocks' in the most recent value of the predictors. Finally, Section 5 concludes the paper.

The Optimal Investment Problem

Following Barberis (2000), the asset allocation framework is characterized by the choice between two assets. The investor allocates *w* to a risky stock portfolio and (1-*w*) to a riskfree asset. With a buy-and-hold investment horizon of K months and preferences given by a constant relative risk aversion utility function, the investor's optimal portfolio allocation maximizes the expected utility of terminal wealth:

$$\max_{\omega} E_T \left[\frac{\left((1 - \omega) \exp\left(r_f K \right) + \omega \exp\left(r_f K + R_{T+K} \right) \right)^{1-\gamma}}{1 - \gamma} \right], \quad (1)$$

where ϒ is the investor's coefficient of relative risk aversion, r_f is the continuously-compounded return on the riskfree asset (assumed constant over the holding period), r_{T+t} is the excess continuously-

compounded return on the stock portfolio in month *T+t*, and the K-period cumulative excess stock return is $R_{T+K} = \sum_{t=1}^{K} r_{T+t}$

The only unknown in (1) is the K-period excess stock return R_{T+K}. To solve the optimization problem, the investor forms an expectation over the distribution of R_{T+K} using a predictive model of stock returns. For expediency in this section, we use general notation for the predictive model and present the specific parameterization in Section 3. Excess stock returns in month T+1 are modelled as a function of the information set z_T and a parameter set θ:

$$r_{T+1} = f(\theta, z_T) + \varepsilon_{T+1}. \qquad (2)$$

The distribution of K-period cumulative excess stock returns follows from iterating (2) forward to the desired forecast horizon. The process of iteration requires an estimate of the parameter set θ and this paper adopts a Bayesian approach. The posterior density of the parameters $p(\theta|z)$ is obtained by combining the investor's prior beliefs about the parameter set *p(θ)* with the likelihood function of the predictive model *L(θ|z)* using Bayes' rule: *p(θ|z)α(θ|z) p(θ)*. The posterior density of model parameters is central to the estimation of the density of future cumulative stock returns R_{T+K}.

Following Barberis (2000), two approaches to obtaining this density are analyzed. First, the predictive model of stock returns (2) is iterated forward using point estimates from $p(\theta|z)$.[5] Treating these point estimates as the true population parameters, the investor forms an expectation of the distribution of future returns conditioned on the fixed parameter estimates $p(R_{T+K}|\hat{\theta}, z)$.[6] Note that this approach to the optimal investment problem fails to account for estimation risk. By merely using point estimates from $p(\theta|z)$, the investor is oblivious to the dispersion of posterior mass around $\hat{\theta}$. This estimation risk, however, is a potentially important source of uncertainty in the asset allocation decision. This is especially true for the predictive regression employed in Section 3 which involves a multivariate system of equations each with several explanatory variables.

The second approach to obtaining the density of future returns allows the investor to incorporate the uncertain nature of parameter estimates, as summarized by the posterior density $p(\theta|z)$. By numerically integrating the predictive likelihood from (2) over the posterior density of the parameters, the density of future returns is conditioned only on data and not on any fixed parameter estimates. We refer to this as the Bayesian predictive density and denote it by *p(R_{T+K}|z)*.

Assessing Return Predictability

Model Parameterization and Relevant Posterior Densities

Our predictive model for excess stock returns uses three predictors (Section 3(ii) details each predictor). Since the time T+1 excess stock return is forecasted using time T values of the predictors, we must also model the dynamics of each predictor in order to generate forecasts of stock returns two or more periods forward. Accordingly, we use a multivariate regression system that captures not only the time-series nature of excess stock return r_t, but also the stochastic evolution of the predictor variables and their contemporaneous relationships:

$$
\begin{bmatrix} r_t \\ DY_t \\ TS_t \\ RB_t \end{bmatrix} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} + \begin{bmatrix} b_{1,1} & b_{1,2} & b_{1,3} \\ b_{2,1} & b_{2,2} & b_{2,3} \\ b_{3,1} & b_{3,2} & b_{3,3} \\ b_{4,1} & b_{4,2} & b_{4,3} \end{bmatrix} \begin{bmatrix} DY_{t-1} \\ TS_{t-1} \\ RB_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t} \\ \varepsilon_{4,t} \end{bmatrix}. \quad (3)
$$

Let $x_t = (DY_t, TS_t, RB_t)$ and $z_t = (r_t, x_t)'$. Given a time-series of T observations, the regression model can be stacked into the form $Z = XC + E$, where $Z = (z_2, ..., z_T)'$, X is the data matrix with rows $(1\ x_1)$, ..., $(1\ x_{T-1})$, and $C \equiv (a, B)$ is the matrix containing regression estimates of the intercepts and slopes of each equation in (3). Along with the covariance matrix of the regression equations Σ, the parameter set to be estimated is $\theta \equiv (C, \Sigma)$.

The posterior density of the parameters $p(\theta|z)$ is estimated using a Markov chain Monte Carlo (MCMC) sampling scheme. Given noninformative priors, Zellner (1971) shows that the conditional/marginal densities of the parameters are:

$$
p\left(\Sigma^{-1}|z\right) \sim \text{Wishart}\left(T - 5, S^{-1}\right)
$$
$$
p(\text{vec}(C)|\Sigma, z) \sim \text{MVN}\left(\text{vec}(\hat{C}), \Sigma \otimes (X'X)^{-1}\right),
$$

where $S = (Z - X\hat{C})'\ (Z - X\hat{C})$ and $\hat{C} = (X'X)^{-1}X'Z$. The MCMC sampling scheme proceeds as follows. We draw an iterate $\Sigma^{-1}[i]]$ from the Wishart density $p(\Sigma^{-1}|z\_\_1jz)$, then conditioning on its inverse $\Sigma^{[i]}$, draw an iterate $C^{[i]}$ from the multivariate Normal density $p(C|\Sigma^{[i]}, z)$. Cycling through this process many times, the iterates $(\Sigma^{[i]}, C^{[i]}$ converge to draw from the joint posterior $p(C, \Sigma\ |z) = p(\theta|z)$. In this paper, numerical estimation of all posterior densities is based on 1,000,000 MCMC iterates.

Having obtained the posterior density of parameters, the distribution of forecasted excess stock returns is estimated. Following Barberis (2000), the multivariate regression is transformed into a vector autoregression. Iterating forward K periods, an expression is obtained for the mean (μ_k) and variance

($\sum_k$) of the K-month cumulative stock return. Given the investment horizon K, the last observed value of the predictors z_T, and the posterior density $p(\theta|z)$, the density of future cumulative stock returns R_{T+K} can be constructed.

As outlined in Section 2, we consider two alternative ways of estimating the distribution of future returns using the posterior density of the parameters. First, estimation risk is ignored by taking the mean values of $p(\theta|z)$ as fixed point estimates $\hat{\theta} = (\hat{C}, \widehat{\sum})$. Inserting these point estimates into the expressions for μ_K and $\sum_K$, a large sample is drawn from the multivariate Normal density with mean μ_K and variance-covariance matrix $\sum_K$. Alternatively, estimation risk surrounding the parameters is incorporated into the predictive density. For each iterate $(C^{[i]}, \sum^{[i]})$ from $p(\theta|z)$, we draw once from the multivariate Normal with the iterate substituted into μ_K and $\sum_K$. This produces the Bayesian predictive density $p(R_{T+K}|z)$ that is conditioned only on data and is independent of θ.

Finally, to control for the effect that the last observed value of the predictors may have in the asset allocation decision, we set z_T in the expression for μ_K equal to the sample mean of the predictor variables. Section 4 explores the impact of the most-recently observed value of z_T more fully.

Data and Results

The data are from the Australian Graduate School of Management (AGSM) monthly stock file from January 1975 to December 1998 (a 288-month time-series). We adopt the AGSM value-weighted market portfolio as our risky stock portfolio. Following Barberis (2000), the riskfree rate used in the asset-allocation decision is set to the December 1998 continuously-compounded return on the 13-week Treasury Note.

Three independent variables are included in the predictive regression. The dividend yield (DY) on the market portfolio is constructed for each month t as the sum of all dividends paid by stocks in the AGSM index from month t-11 to the end of month t divided by the price level of the stocks at the end of month t. The term spread (TS) is the difference between yields on 10-year Australian Government Bonds and 13-week Treasury Notes. Finally, the relative bill rate (RB) is the annualized return on the 13-week Treasury Note less its 12-month moving average.

The point estimates for $\theta \equiv (a, B, \sum)$ are shown in Table 1. The top panel reports the mean of the posterior densities for the regression intercepts *(a)* and slope coefficients (B), with the standard deviation of the MCMC iterates in parentheses. The bottom panel reports variances from the diagonal of $\sum$, while off-diagonal elements have been converted to correlations. Examining the top row of B, one-step-ahead excess stock returns are positively related to dividend yield (0.8948), and negatively related

to the term spread (-0.1439) and the relative bill rate (-0.2570). If we were to compare the point estimates to their standard deviations and conduct a classical test of the hypothesis that each estimate equals zero, we are likely to conclude that only dividend yield has predictive ability. In contrast, Section 4 documents the extent to which each predictor impacts on an investor's optimal asset allocation decision.

The second point to note from Table 1 is that all the predictors are highly persistent but mean-reverting. This is evident from the autoregressive coefficients of 0.9191 for the dividend yield, 0.9151 for the term spread, and 0.9039 for the relative bill rate. The mean-reverting nature of the predictor variables has important implications for the analysis in Section 4.[7] Finally, examining the correlation of the predictors with contemporaneous excess stock returns indicates that the dividend yield works in the opposite direction to the relative bill rate and the term spread. The former induces negative serial correlation in excess stock returns – a positive shock to the dividend yield forecasts an increase in future excess returns, while the negative correlation is related to a decrease in current excess stock returns. The latter would appear to induce positive serial correlation in excess stock returns – positive shocks to the term spread and the relative bill rate forecast lower future returns while also being related to lower current excess stock returns. Although simplified because it negates the correlation between the predictor variables, it remains to be seen which of these effects dominates the investor's asset allocation decision.

In assessing the economic significance of return predictability, we adopt as our benchmark the optimal allocation of an investor who assumes returns are unpredictable (i.e. an *i.i.d*. return model is utilized to forecast the distribution of K-period cumulative returns). In this case, Barberis (2000) shows numerically that the optimal allocation is independent of the investment horizon K. Thus, the extent to which the optimal allocation using the predictive regression (3) differs from the constant allocation indicates the economic significance of return predictability.

Table 1

Parameter Estimates from the Multivariate Regression Model

a	B		
−0.0346	0.8948	−0.1439	−0.2570
(0.0165)	(0.3576)	(0.1716)	(0.3617)
0.0035	0.9191	0.0067	0.0269
(0.0008)	(0.0183)	(0.0088)	(0.0186)
−0.0005	0.0274	0.9151	−0.0671
(0.0023)	(0.0498)	(0.0239)	(0.0504)
0.0022	−0.0505	−0.0040	0.9039
(0.0012)	(0.0257)	(0.0123)	(0.0259)

Σ			
0.0030	−0.6989	−0.0025	−0.3282
(2.5841e−4)	(0.0306)	(0.0597)	(0.0533)
	7.9775e−6	0.0286	0.2184
	(6.7923e−7)	(0.0596)	(0.0568)
		5.8872e−5	0.1883
		(5.0039e−6)	(0.0575)
			1.5632e−5
			(1.3316e−6)

Notes:
Results in this table are for the multivariate predictive regression (3). The reported numbers in the top panel are the mean and standard deviation (in parentheses) for the posterior distributions of regression intercepts (a) and slope coefficients (B). The bottom panel reports variances from the diagonal of Σ, while off-diagonal elements have been converted to correlations. Results cover the period January 1975 to December 1998 and are based on 1,000,000 MCMC iterates.

Figure 1 reports the optimal portfolio allocation of an investor using the dividend yield, term spread, and relative bill rate to predict excess stock returns, both with and without estimation risk. To provide a benchmark (dash-dot lines), the constant allocation of an unpredictable return investor with Y=5 (10) is 46% (23%) in the risky stocks. Ignoring estimation risk (solid line), when an investor uses the three variables to predict excess stock returns, the optimal allocation to risky stocks is a positive function of her investment horizon. For example, an investor with Y=5 and a 10-year investment horizon allocates 55% to stocks. This result indicates that the negative serial correlation induced by the dividend yield dominates the effect of the term spread and the relative bill rate in the asset allocation decision. Negative serial correlation in forecasted stock returns slows the evolution of the conditional variance of cumulative excess stock returns. As a result, the standard deviation of the distribution of future returns at a 10-year investment horizon is 45% for the predictable return investor compared to 60% for the unpredictable return investor, leading to a higher allocation to risky stocks when the predictive model is employed.

Figure 1
Optimal Asset Allocations: Predictable and Unpredictable Stock Returns
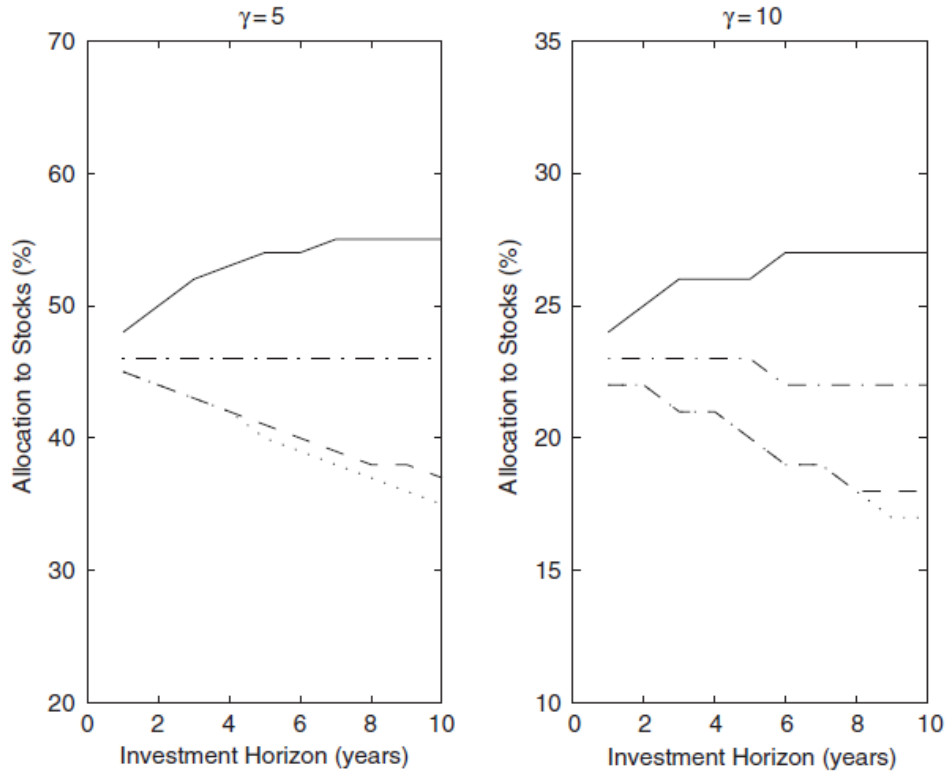
Notes:

This figure presents optional asset allocation plots when the investor believes excess stock returns follow the multivariate prediction model (3). The solid line indicates the investor ignores estimation risk by using the distribution of future returns conditional on fixed parameter estimates. The dotted line indicates the investor uses the Bayesian predictive density to incorporate estimation risk. The dash-dot (dash) lines are the optimal allocation from the *i.i.d.* model of returns without (with) incorporating estimation risk. Left and right plots are for an investor with $\gamma = 5$ and 10 respectively.

An interesting comparison can be made between the results in Figure 1 and those for US data reported by Barberis (2000). The horizon effects using Australian data are notably smaller. While Australian return predictability induces an investor with Υ=5 and a 10-year horizon to allocate 9% more to stocks, her US counterpart allocates 30% more to stocks. There are two possible explanations for this result. First, Barberis (2000) uses only the dividend yield to predict, while this paper uses three predictors in the asset allocation decision. Results (not reported) which replicate Barberis (2000) by using dividend yield as the sole predictor produce horizon effects virtually identical in magnitude to those reported in Figure 1. The more likely explanation is that Australian correlation between current excess stock returns and dividend yield is much smaller than in Barberis (2000). This has the effect of

decreasing the magnitude of negative serial correlation in the excess stock return series, leading to a higher forecast of the variance of cumulative returns, and thus reducing the horizon effects.

When estimation risk surrounding the estimates from the predictive model is incorporated (dotted line), the results are very different – the investor allocates less to stocks at longer horizons (35%). This finding differs considerably from Barberis (2000), who reports that, while estimation risk causes the investor to reduce her allocation to stocks at longer investment horizons, the allocation for predictable returns remains well above the allocation for the *i.i.d*. returns case.

At first glance, negative horizon effects may appear counterintuitive. Barberis (2000) argues that, if returns are predictable, the variance of long-horizon returns grows slower than linearly with investment horizon, making stocks look relatively less risky at longer horizons leading to higher allocations. However, there are many parameters to estimate in the multivariate prediction model (3), introducing a considerable degree of estimation risk. When accommodated in the asset-allocation decision, the perceived riskiness of the stock investment grows faster than linearly with the investment horizon and the optimal allocation to stocks decreases. In addition, the negative horizon effect may be exacerbated by the modelling of buy-and-hold investments. The consequences of incorrectly assessing the degree of predictability are more severe when the investor is locked-in for long horizons.

In summary, examining the horizon effects induced by return predictability produces two main results. When an investor ignores estimation risk, the degree of predictability in Australian returns appears sufficient to cause the investor to allocate more heavily to stocks at longer investment horizons. These horizons effects, however, are much smaller than reported for US data. When the additional uncertainty induced by estimation risk is incorporated into the asset allocation decision, the evidence of predictability disappears. This finding contrasts sharply with Barberis (2000), who reports that while estimation risk diminishes the horizon effects, the predictability is sufficiently strong to leave some positive effects. The uncertainty over the true parameters of the predictive model seems so severe that the investor allocates about the same as under the *i.i.d* case.[8] Using this metric, it would appear that return predictability has little impact on the investor's optimal investment decision. This conclusion should be interpreted with caution, however, given that the analysis to date has set the predictor variables z_T at their sample means. This issue is explored further in the next section.

Assessing the Importance of Each Predictor

The previous section assessed the economic significance of return predictability by examining the horizon effects induced by all three variables jointly. This section measures the relative importance

of each predictor. Specifically, the impact of each predictor variable on the asset allocation decision is documented.

Recall that we have a predictive model of excess stock returns using three predictors – dividend yield on the market portfolio, the term spread and the relative bill rate. In the classical framework, it is common to run a series of regressions involving combinations of the predictors and judge relative importance using t-statistics and R^2's. Table 2 presents such an analysis. This approach is strictly statistical in nature. The economic implications of the predictors, however, are somewhat more ambiguous. For example, the t-statistics on term spread (model b) and relative bill rate (model c) are insignificant; does this mean these variables are of no value to investors? Similarly, moving from model (a) to model (d) by including term spread increases R^2 by 0.0022 (but decreases adjusted R2 by 0.0013); does this mean term spread is useful to investors? Finally, while dividend yield appears to have some predictive ability, would a buy-and-hold investor with a 1-year investment horizon assign the same degree of importance to dividend yield as an investor with a 10-year horizon? Statistical analysis does not answer these questions.

The asset allocation framework provides an ideal setting to answer such questions. Irrespective of its statistical importance, a predictor is economically important if its observed value affects an investor's optimal asset allocation. Our approach follows Kandel and Stambaugh (1996) who measure the sensitivity of the investor's optimal portfolio to the most-recently observed value of each predictor. Intuitively, if a variable cannot predict excess stock returns, the optimal allocation will be insensitive to its last observed value.

The optimal asset allocations in Section 3 were derived by setting the last observed value of the predictors $x_T=(DY_T, TS_T, RB_T)$ equal to their sample mean over the January 1975 through December 1998 estimation period. Current values of the predictors x_T enter the allocation decision through the predictive density of future returns R_{T+K}, since the mean of this density μ_K is a function of $z_T=(r_T, x_T)$. Hence, changing x_T affects the mean of the distribution of future returns.

To assess the sensitivity of the optimal portfolio to the last observed value of each predictor, we 'shock' one predictor variable (one or two standard deviations above or below its mean), while holding the other predictors at their means. In essence, we measure the investor's response to hypothetical values of the predictors. The divergence between the optimal portfolio conditioned on the shocked value of the predictor and the optimal portfolio when all predictors are set at their mean highlights how economically important that variable is to the investor.

Figure 2
Sensitivity of Optimal Asset Allocation to Shocks to Predictors
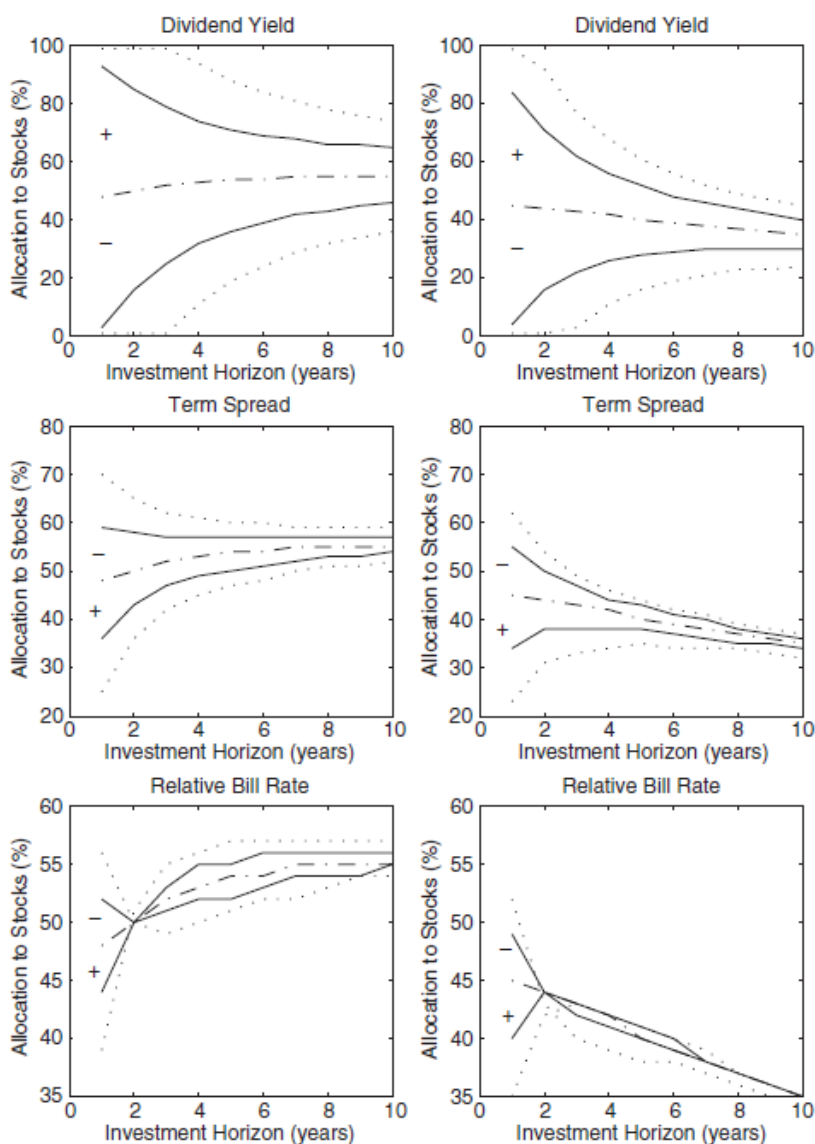
Figure 2 displays the optimal portfolio allocations when the dividend yield, term spread, and relative bill rate are set to ∓1 or ∓2 standard deviations from their sample mean, while holding the other predictors at their sample means. Solid (dotted) lines indicate that the predictor has been set one (two) standard deviations from its sample mean, with the sign of the shock indicated by the plus and

minus signs. Dash-dot lines indicate that all predictors are at their sample means (exactly the same results shown in Section 3). In the left (right) panels, the investor ignores (incorporates) estimation risk into the allocation decision. The investor has Υ=5 in all diagrams.

The most striking result from Figure 2 is that, at short to medium investment horizons, the optimal allocation to stocks is highly sensitive to the most-recent value of each predictor. This is most evident for the dividend yield (top panels), where a two standard deviation increase (decrease) causes the investor to allocate all (none) of her wealth to stocks for horizons of up to three years. Likewise, shocks to term spread (middle panels) impact on allocations out to at least the medium term. Shocks to the relative bill rate (bottom panels) only affect asset allocations over very short investment horizons.

Values of the predictors two standard deviations from their means are not unrealistic. Although uncommon, such shocks do fall within the maximum and minimum values observed during the 288-month sample (summary statistics are not reported). Regardless, the sensitivity of allocations remains evident for shocks of one standard deviation, although the magnitude is reduced.

With respect to distinctions between statistical and economic significance, shocks to term spread (middle panels) are particularly interesting. In Table 2, the t-statistic on term spread is clearly insignificant, possibly leading a classicist to conclude that the term spread has no predictive ability and therefore is of no value to investors. In stark contrast, Figure 2 reveals that when term spread is two standard deviations above (below) its sample mean, an investor will allocate 25% (70%) to equities at a 1-year investment horizon, compared to 48% when term spread is set at its sample mean. Clearly, knowledge of the term spread is relevant information to short-horizon investors.

Another feature of these figures is the insensitivity of the optimal allocation to initial values of the predictor variables at the 10-year horizon (especially when estimation risk is incorporated). As investment horizon increases, the lines tend to converge suggesting that a current shock to a predictor is relatively unimportant for long-horizon investors. This result is appealing because it indicates that the predictor variables come from stationary times-series processes – shocks to the times-series do not have permanent effects. This result differs from Barberis (2000) who finds that shocks to the dividend yield imply investment allocations that do not converge at long investment horizons.

One set of results that is difficult to explain is the sensitivity of asset allocations to the relative bill rate in Figure 2 (bottom panels). Given the negative relationship between the relative bill rate and one-step-ahead excess stock returns (see Table 1), the direction of these plots is as expected – an increase (decrease) in the relative bill rate decreases (increases) the mean of the distribution of future returns causing the investor to allocate less (more) heavily to the risky asset. The puzzling result,

however, is that the optimal allocation following a decrease in the relative bill rate is a non-monotonic function of the investment horizon – after the initial shock, the optimal portfolio initially decreases then increases with investment horizon. Although this result is difficult to explain, an examination of the bottom-right plot in Figure 2 indicates that the optimal allocation for an investor who accounts for estimation risk converges for investment horizons beyond two years. The main conclusion we can draw is that the investor's optimal allocation is somewhat sensitive to the relative bill rate, but this sensitivity is very short lived.

The final result to come from Figure 2 is that, although estimation risk dampens the impact of a shock to a predictor, the investor's optimal allocation is still sensitive to the most recently observed value of each predictor (refer to the right side plots). This result is quite intuitive – when the investor incorporates estimation risk into her allocation decision, they are acknowledging the uncertain nature of the predictive evidence. If the investor is less certain about the predictive ability of the variables, the optimal allocation will be less sensitive to the last observed values of the variables.

Kandel and Stambaugh's (1996) metric for examining the economic significance of return predictability also serves as a useful way of thinking about the relative importance of each predictor to an investor. Examining the change in the optimal portfolio for a standard change in each predictor gives an indication of the relative weight the investor places on the predictive ability of each variable. The larger the change in the optimal portfolio, the more importance the investor assigns to that variable's predictive ability.

Figure 3 represents in a single plot the optimal asset allocations for shocks of positive one standard deviation to each variable separately while holding the other predictor variables at their sample means. The dotted, dash-dot, and dashed lines indicate shocks to the dividend yield, term spread, and relative bill rate respectively. The solid line indicates the optimal allocation when all predictors are set at their sample means. The left (right) panel ignores (incorporates) estimation risk. Again, the investor has $\Upsilon=5$.

The clear result to come from Figure 3 is that the dividend yield has the largest impact on the asset allocation decision. This coincides with the statistical evidence on the issue, with the dividend yield being the only statistically significant variable in the predictive regression (model g). To conclude that the statistical approach and the asset allocation approach yield identical results is, however, premature. Using statistical analysis, the term spread and the relative bill rate would both be considered insignificantly different from zero and hence, not of importance to investors. Within an asset allocation

framework, Figure 3 shows that the term spread has far greater impact than the statistical evidence suggests.

Figure 3
Relative Impact of Each Predictor Variable
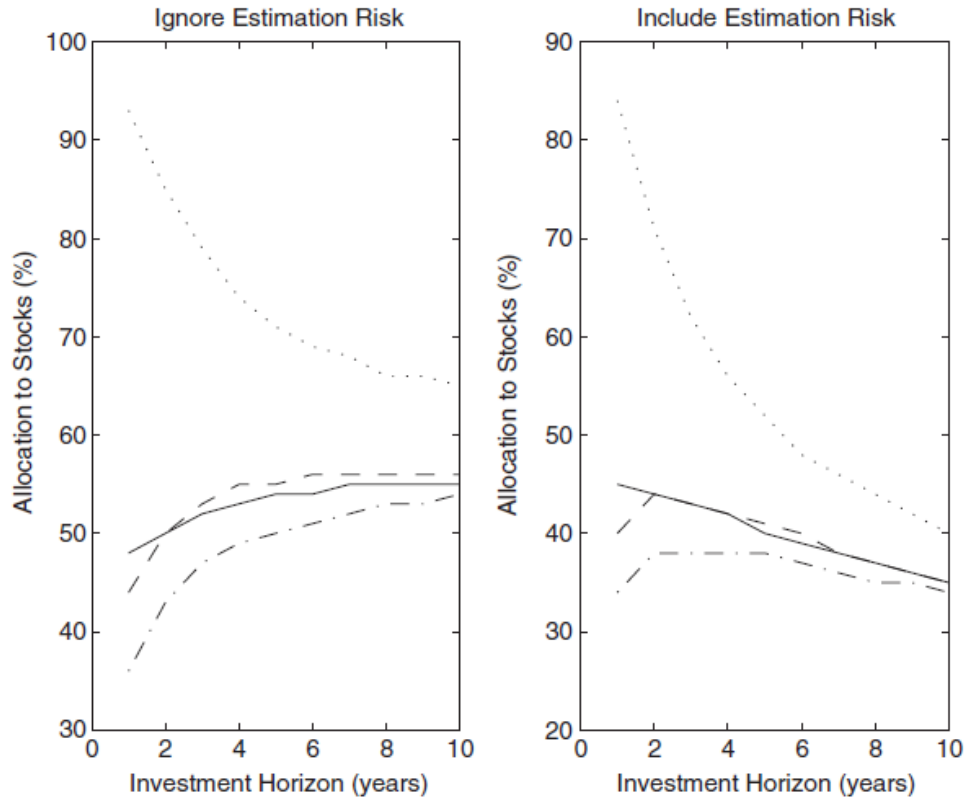
Notes:

This figure shows the relative impact on optimal asset allocations from a shock to the most recently observed value of a predictor. Dotted, dash-dot, and dashed lines indicate that the dividend yield, term spread, and relative bill rate have been set one standard deviation above their sample means respectively, while holding the other predictor variables at their sample means. Solid lines indicate that all predictor variables have been set at their sample means. The left (right) plot is for an investor ignores (includes) estimation risk into the asset allocation decision. The investor has coefficient of risk aversion $\gamma = 5$.

Conclusions

This paper draws a distinction between the statistical significance of return predictability and the influence that predictor variables have on a risk-averse investor's actions. Classical statistical inference does not necessarily capture the economic importance with which investors regard various predictors. Instead, we use a unified framework for portfolio choice under uncertainty to examine the

extent to which investors value the evidence of return predictability. If a variable is useful in predicting stock returns, knowledge of that variable will have an impact on the portfolio allocation of a utility-maximizing investor. Thus, the evidence on return predictability is judged from an economic perspective, not a statistical one.

This paper builds on prior asset allocation studies in several directions. Our predictive model is a multivariate system of equations involving one risky asset and three predictors. The extent of predictability is judged by the horizon effects induced by the three variables jointly. We are also able to assess the relative importance of each predictor by the sensitivity of optimal asset allocations to shocks in the most-recent value of each predictor in turn. While studies such as Avramov (2002) have included multiple predictors into an asset allocation decision, the relative impact of each predictor is not measured. Finally, we utilize a dataset which has not been used before in the asset allocation literature. Some notable differences between our results and those of Barberis (2000) are documented, leading us to caution against generalizing findings on predictability too broadly.

Our results can be summarized as follows. First, classical statistical inference can belie the economic importance of predictor variables. The term spread and, to a lesser extent, the relative bill rate show no sign of statistical significance, yet optimal portfolio allocations are sensitive to their most-recently observed values. Dividend yield has the largest impact on allocation decisions, a fact also detected by classical inference. Second, estimation risk is an important component of the total uncertainty facing an investor. Ignoring estimation risk, a risk-averse investor using the multivariate model to predict stock returns allocates around 9% more to stocks at an investment horizon of 10 years as compared to a 1-year investment horizon. When estimation risk is taken into account, the horizon effects induced by predictability disappear. In fact, the investor allocates around 10% less to stocks at the 10-year investment horizon.

Reporting results for US data, Barberis (2000) finds much larger horizon effects from return predictability, despite the fact that dividend yield is the sole predictor. At a 10-year investment horizon, a US investor allocates 30% more to stocks perhaps indicating that US returns are more predictable than Australian returns. Unlike the current findings, Barberis reports that the incorporation of estimation risk only reduces the magnitude of the horizon effects. This difference may be due to the different sample sizes – Barberis estimates a predictive model using 44 years of data; our estimation uses only 25 years. It follows logically that our model is estimated with less precision, introducing a higher degree of estimation risk. Issues such as these are likely to be relevant if the economic significance metric is adopted to assess return predictability in other countries.

There are several avenues for further research on this topic. Perhaps the most important is the incorporation of periodic portfolio rebalancing into the allocation framework. With dynamic rebalancing, portfolio allocations will incorporate a component that hedges against changes in the opportunity set. The asset allocation framework could also be generalized to allow the choice between multiple stochastic assets, rather than just the market portfolio and a (constant) riskfree asset. For example, adding an international stock portfolio would allow the investor to diversify and hedge internationally.

References

Avramov, D. (2002), 'Stock Return Predictability and Model Uncertainty', *Journal of Financial Economics*, Vol. 64, pp. 423–58.

Barberis, N. (2000), 'Investing for the Long Run when Returns are Predictable', *Journal of Finance*, Vol. 55, pp. 225–64.

Bawa, V.S., S.J. Brown and R.W. Klein (1979), *Estimation Risk and Optimal Portfolio Choice* (1st ed., Amsterdam: North Holland).

Brennan, M.J., E.S. Schwartz and R. Lagnado (1997), 'Strategic Asset Allocation', *Journal of Economic Dynamics and Control,* Vol. 21, pp. 1377–403.

Campbell, J.Y. (1987), 'Stock Returns and the Term Structure', *Journal of Financial Economics*, Vol. 18, pp. 373–99.

———— (1991), 'A Variance Decomposition for Stock Returns', *The Economic Journal*, Vol. 101, pp. 157–79.

———— and L.M. Viceira (1999), 'Consumption and Portfolio Decisions when Expected Returns are Time Varying', Quarterly *Journal of Economics,* Vol. 114, pp. 433–95.

———— ———— (2001), 'Who Should Buy Long-term Bonds?', *American Economic Review*, Vol. 91, pp. 1–29.

Cochrane, J.H. (1999), 'New Facts in Finance', *Economic Perspectives Federal Reserve Bank of Chicago,* Vol. 23, pp. 36–58.

Fama, E.F. and K.R. French (1988), 'Dividend Yields and Expected Stock Returns', *Journal of Financial Economics,* Vol. 22, pp. 3–25.

———— ———— (1989), 'Business Conditions and Expected Returns on Stocks and Bonds', *Journal of Financial Economics,* Vol. 25, pp. 23–49.

Hodrick, R.J. (1992), 'Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement', *Review of Financial Studies,* Vol. 5, pp. 357–86.

Kandel, S. and R.F. Stambaugh (1996), 'On the Predictability of Stock Returns: An Asset-allocation Perspective', *Journal of Finance*, Vol. 51, pp. 385–424.

Keim, D.B. and R.F. Stambaugh (1986), 'Predicting Returns in the Stock and Bond Market', *Journal of Financial Economics,* Vol. 17, pp. 357–90.

Klein, R.W. and V.S. Bawa (1976), 'The Effect of Estimation Risk on Optimal Portfolio Choice', *Journal of Financial Economics*, Vol. 3, pp. 215–31.

Lynch, A.W. (2001), 'Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability', *Journal of Financial Economics*, Vol. 62, pp. 67–130.

Zellner, A. (1971), *An Introduction to Bayesian Inference in Econometrics* (1st ed., New York: John Wiley and Sons).

---

[1] See, for example, Keim and Stambaugh (1986), Campbell (1987 and 1991), Fama and French (1988 and 1989) and Hodrick (1992).

[2] See, for example, Barberis (2000), Brennan et al. (1997), Campbell and Viceira (1999 and 2001) and Lynch (2001).

[3] Bawa, Brown and Klein (1979) provide a comprehensive discussion of the importance of recognizing estimation risk in a variety of decision contexts.

[4] See Fama and French (1989), Campbell (1991) and Hodrick (1992).

[5] Where point estimates are required, we take the mean of the posterior density $p(\theta|z)$ and denote it $\hat{\theta}$.

[6] This is very similar in nature to the maximum likelihood approach taken in asset allocation papers such as Lynch (2001).

[7] Note that the equations in (3) which model the dynamics of the predictors have lagged values of all three variables. While we could impose a prior restriction on the B matrix that each predictor is a first-order autoregressive process (i.e. off-diagonal coefficients are zero), Table 1 shows that the off-diagonal coefficients are close to zero anyway.

[8] At the 10-year horizon, the allocation to stocks after incorporating estimation risk (dotted line) is almost identical for the predictive model (35%) and the i.i.d model (37% dashed line).