Consumer Expectations and Short Horizon Return Predictability*

Egon Kalotay
Department of Accounting and Finance
Macquarie University
ekalotay@efs.mq.edu.au

Phil Gray
Business School
University of Queensland

Samantha Sin
Department of Accounting and Finance
Macquarie University

Working draft: May 30, 2005

*This version is a draft. Please do not quote without permission. Comments and suggestions welcome.

Abstract

We study the role of consumer expectations, as measured by consumer behavior (departures from an intertemporal budget constraint) and the University of Michigan's Index of Consumer Sentiment, in modeling time variation of expected equity returns over short horizons. We find strong evidence of return predictability based on estimates of short term departures from the budget constraint, and the predictability is evident after accounting for various sources of estimation risk. However, the apparent predictability does not necessarily give rise to useful signals based on shifts in aggregate consumption and and the components of aggregate wealth due to estimation risk associated with the budget constraint. We find that the survey based measure of expectations complements the behavioral measure but has no apparent stand-alone predictive value in forecasting equity returns.

1 Introduction

A degree of systematic, unexploited equity return predictability is consistent with modern asset pricing theory. Evidence of return predictability is not necessarily evidence of an anomaly, market inefficiency or investor irrationality. However, there is little apparent consensus of opinion on the degree of return predictability and its economic implications - though the evidence favoring some degree of predictability - particularly over long horizons - is arguably gaining wider acceptance.[1]

A recent study by Lettau and Ludvigson (2001) (henceforth LL) breaks with the tradition of most previous studies of predictability in two important ways. First, while existing evidence of predictability using 'known' predictors (such as dividend yield) tends to arise at longer horizons, LL provide empirical evidence of predictability over short (quarterly) horizons. Second, LL's predictive variable has sound *ex-ante* theoretical foundations, suggesting strong links between variation in consumer expectations, as captured by the predictive variable $\widehat{cay}$, and subsequent equity returns.

LL model a long-run dynamic equilibrium relation between the log of consumption growth (c), asset returns (a) and labour income (y), with deviations from the cointegrating relation named cay. Fluctuations in cay can be shown to reflect expectations of return on the true market portfolio in a general, multiperiod consumption-based framework. While cay is not directly observable, LL construct an empirical estimate - denoted by $\widehat{cay}$ - and argue that the important predictive components of cay are captured by $\widehat{cay}$. As such, $\widehat{cay}$ can be expected to forecast equity returns by virtue of the Granger Representation Theorem (GRT). Their findings suggest that $\widehat{cay}$ does indeed have significant ability to forecast the quarterly market returns in-sample.

Such findings are controversial and subsequent works - such as Brennan and Xia (2002), Hahn and Lee (2002) and Rudd and Whelan (2002) - have critiqued the LL methodology and questioned their assumptions and findings. This study presents empirical evidence related to two particular avenues of criticism.

First, we examine whether findings of predictability are robust to the assumption that cointegrating parameters are fixed and known. Whilst Lettau and Ludvigson argue this assumption is reasonable - given that it's an implication of the theory, and given the convergence properties of parameter estimates if the series are cointegrated - their argument rests on a several assumptions and approximations of un-tested empirical validity. Our results suggest that, conditional on model specification, the economic significance of short-horizon return predictability is not driven by the assumption that cointegrating parameters are fixed and known - even though we do find substantial uncertainty associated with the parameters of cointegration.[2]

Second, we examine whether an alternative survey-based measure of consumer expectations supplants or complements $\widehat{cay}$ as a predictor of return. Measures of consumer confidence, such as the Index of Consumer Sentiment constructed by the University of Michigan, represent an alternative source of information about aggregate expectations of economic fundamentals underlying consumption and investment choices. Further, in the presence of 'strategic complementarities' the sentiments themselves may drive economic outcomes. The information Lettau and Ludvigson suggest is embodied in $\widehat{cay}$ may be accessible in survey based measures of sentiment - without making the unreasonable assumptions about consumers' decision processes noted by Brennan and Xia (2002).

We examine the economic implications of our findings in the context of a simple, stylized asset allocation decision, and by studying the long horizon predictive properties of predictive parameters estimated over quarterly horizons. In so doing we find that evidence of predictability does not necessarily mean that shifts in the aggregate consumption, asset wealth and income provide economically significant signals about expected returns to a risk averse investor who does not have non-sample information about the cointegration parameters.

To re-emphasize LL's main finding, consider regressing quarterly S&P500 excess index returns (r_t) on lagged $\widehat{cay}$ over the period 1951 Q4 through 2003 Q1. The estimated predictive regression ($t-$statistics in parentheses) using this extension of LL's sample is:

$$
\begin{array}{rcl}
r_t & = & \underset{(3.23)}{0.018} + \underset{(4.60)}{2.0494}\, \widehat{cay}_{t-1} + e_t \qquad (R^2 = 0.0939)
\end{array}
$$

The most striking finding is the in-sample predictability of $\widehat{cay}$ as summarized by the regression R^2 of 9.39%. This level of short-horizon predictability is significantly higher than that which obtains using common alternatives. For example, if we substitute the lagged dividend yield (DY) for $\widehat{cay}$, the regression R^2 over the corresponding sample is a mere 1.86%.[3]

The apparently strong predictive properties of $\widehat{cay}$ have been questioned by researchers on several fronts. First, Brennan and Xia (2002) and Hahn and Lee (2002) question the use of a predictor constructed using in-sample data. That is, the parameters used to construct $\widehat{cay}$ are estimated over the period used to test predictability. Second, Brennan and Xia (2002) question the underlying behavioral assumptions. Specifically, cointegration of c, a and y requires that the representative agent whose consumption behavior is captured by $\hat{cay}$ is able to solve an extremely difficult dynamic consumption-investment plan so that his current consumption-wealth ratio depends on his expectations of future risky investment opportunities. Third, Rudd and Whelan (2002) question the validity of the empirical assumptions underlying construction of $\hat{cay}$.

Lettau and Ludvigson (2002) respond to Brennan and Xia (2002) and Hahn and Lee (2002) with reference to economic theory and the GRT. First, they note that knowledge of $\hat{cay}$ parameters (wealth shares) is an implication of the theory and argue that out of sample testing does not safe-guard against spurious findings of in-sample predictability. Second, if if c, a, and y are cointegrated, the GRT implies that lagged $\hat{cay}$ must forecast growth in the stock-market component of a, given that it has no apparent ability to forecast, c, y or non-stock market wealth.[4]

The current paper contributes to the debate over the predictive power of *cay* by extending the LL analysis in several new directions. First, we assess the economic importance of return predictability by examining the impact of *cay* on simple asset allocation choices. There is a growing literature that emphasizes the importance of studying predictability in terms of its impact on economic choices - not just statistical significance. Like Kandel and Stambaugh (1996) and Barberis (2000), we utilize a Bayesian econometric approach. We formally account for estimation risk in both the predictive regression *and* the cointegrating model used to construct $\hat{cay}$. In doing so, the methodology proposed in the current paper mitigates some of the concerns raised by Brennan and Xia (2002) and Hahn and Lee (2002). By adopting a Bayesian perspective we need not treat the parameters governing the joint dynamics of c, a and y as fixed and known. Rather, we characterize and account for the uncertainty in modeling *cay* and conditional expectations of the return premium based on *cay*. This is the second general contribution of the paper.

Finally, we propose and investigate the use of a consumer sentiment measure as a predictor of short-horizon equity returns. LL's consumption-to-wealth ratio can be thought of as a behavioral summary of expectations about future investment opportunities assuming the representative, rational agent solves a difficult dynamic consumption-investment plan. Whether the plausibility of this assumption affects the credibility of tests of its implications is an example of a more fundamental debate over whether rationality can be defined in terms of the choices it produces or the process that is used to make the choices.[5] A pragmatic response to this line of questioning (in this particular instance) is to shift focus from the behavioral implications of the model to a direct measure of the expectations that influence agents' behavior.

In the consumption based framework underlying the use of $\hat{cay}_{t-1}$ as a measure of expectations, departures from the steady state ratio of consumption to wealth occur during times of optimism (characterized by higher consumption and lower risk aversion) and pessimism (lower consumption and higher risk aversion). An alternative to measuring departures of consumption from its steady state relationship with asset wealth is to use a measure that attempts to quantify agents optimism or pessimism directly - without explicit reference to optimizing behavior based on an intertemporal budget constraint.

The Index of Consumer Sentiment (ICS), devised in the 1940's by George Katona and published since 1966 by the University of Michigan, provides a composite measure of consumers' perceptions of their 'ability and willingness to buy'. Specifically, the index seeks to capture consumers' ability and willingness to undertake discretionary expenditures.[6]

To the extent that the optimism or pessimism captured by the ICS is an alternative measure of the expectations driving fluctuations in $\hat{cay}_t$ in a rational consumption based framework, the ICS should also forecast equity returns. Further, if fluctuations in $\hat{cay}_t$ are driven by factors other than changes in consumer expectations of asset returns, the ICS may complement, supplant or augment $\hat{cay}_t$ as a predictor of return.

Alternatively, returns may be predicted by shifts in consumers sentiment because the expectations are self-fulfilling. We need not appeal to the consumption based framework if we admit the possibility of *strategic complementarities* or *coordination failures*. In these circumstances, observed economic outcomes can reflect expectations themselves rather than economic fundamentals that give rise to (or presage) such outcomes. At a macro-level, Matsusaka and Sbordone (1995) provide evidence consistent with the existence of strategic complementarities using the ICS as a measure of expectations. Upon controlling for a wide range of economic conditions (fundamentals), they consistently reject the null hypothesis that consumer sentiment does not Ganger cause GNP.

We argue that the ICS may provide an alternative (or complementary) source of information about the expectations captured by $\hat{cay}$. As such, we examine whether this survey-based measure of consumer sentiment is of stand alone or incremental value in forecasting equity returns.

2 Expected Returns and $\hat{cay}$: A Review of the Argument

This section provides a brief overview of the theoretical framework linking consumption, aggregate wealth and expected returns, as well as describing how the predictive components of the consumption-to-wealth ratio can be expressed in terms of observables. The reader is referred to LL for a full exposition.

In a representative agent economy, agents' intertemporal budget constraint implies:

$$c_t - w_t = E_t \sum_{i=1}^{\infty} \rho_w^i (r_{w,t+i} - \Delta c_{t+i}), \tag{1}$$

where c is log of aggregate consumption, w is log of aggregate wealth, r_w is the return on the portfolio of aggregate wealth, and ρ_w is the steady-state ratio of new investment to aggregate wealth. The budget constraint (1) implies that changes in the consumption-to-wealth ratio $(c_t - w_t)$ forecast market returns $(r_{w,t+i})$ or changing consumption growth (Δc_{t+i}).

The consumption-to-wealth ratio, however, cannot be used directly to forecast. Aggregate wealth (W_t) comprises asset holdings (A_t) and human capital (H_t), the latter being unobservable.

To circumvent this problem, LL assume that the non-stationary components of human capital H_t are well described by aggregate labour income Y_t, such that:

$$h_t = \kappa + y_t + z_t, \tag{2}$$

where z_t is a mean-zero stationary random variable and lower case variables denote logs. LL suggest a number of possible rationalisations for the approximation (2).

Consequently, aggregate wealth can be approximated by:

$$w_t \approx \xi a_t + (1 - \xi) h_t, \tag{3}$$

with ξ equal to the mean of asset wealth as a proportion of total wealth. The return on aggregate wealth is decomposed accordingly:

$$r_{w,t} \approx \xi r_{a,t} + (1 - \xi) r_{h,t}, \tag{4}$$

where $r_{a,t}$ and $r_{h,t}$ are time t log returns on asset wealth and human capital respectively. Substituting (2) and (4) into (1) yields the following re-statement of the (log) consumption-to-wealth ratio:

$$\begin{aligned}\hat{cay}_t &\equiv c_t - \xi a_t - (1-\xi)y_t \\ &= E_t \sum_{i=1}^{\infty} \rho_w^i \left\{[\xi r_{a,t+i} + (1-\xi) r_{h,t+i}] - \Delta c_{t+i}\right\} + (1-\xi) z_t. \end{aligned} \tag{5}$$

Equation (5) is the consumption-to-wealth ratio expressed in terms of observables – a linear combination of c_t, a_t and y_t Lettau and Ludvigson refer to as $\hat{cay}_t$. From (5), it is clear that the residual $\hat{cay}_t$ reflects expectations of returns on asset wealth ($r_{a,t+i}$) as long as the returns on human capital ($r_{h,t+i}$) and consumption growth (Δc_{t+i}) are relatively stable (or as LL note, highly correlated with $r_{a,t+i}$). Importantly, if each component of the right-hand side of (5) is stationary and each of c_t, a_t and y_t are unit root processes, then log consumption, aggregate wealth and labour income are cointegrated, and $\hat{cay}_t$ is a cointegrating residual. Henceforth, $\hat{cay}$ or $\hat{cay}_t$ will be used to denote an empirical point estimate of the cointegrating residual. Henceforth, unless stated otherwise, $\hat{cay}$ refers to $\hat{cay}_{t-1}$.

3 Econometric Approach

Two steps are required to use $\hat{cay}$ in a predictive regression. First, to compute trend deviations ($\hat{cay}$), the cointegrating parameters must be estimated. LL follow Stock and Watson (1993) in using dynamic least squares to obtain consistent estimates of the cointegrating parmaters. Second, the time-series of $\hat{cay}$ can then be used as the independent variable in the predictive regression.

In this paper we adopt a Bayesian approach and thereby estimate posterior densities of the parameters of both the predictive and cointegrating regressions to account for estimation risk in measures of return predictability. Kandel and Stambaugh (1996) and Barberis (2000) show that conclusions regarding the degree of predictability can change significantly when estimation risk is taken into account. An added advantage of the Bayesian approach is that, rather than assuming knowledge of fixed wealth shares, the posterior distribution of the cointegrating parameters can be utilized to assess the evidence of predictability conditional on c, a and y only (as opposed to conditioning on $\hat{cay}$ and ignoring the uncertainty surrounding parameter estimates in the cointegrating regression). This goes some way towards alleviating the concerns of Brennan and Xia (2002) over LL's in-sample estimate of $\hat{cay}$ that's assumed fixed and known.

In the following section, we characterize the posterior distributions of the cointegrating weights and the adjustment coefficients, conditional on model specification.

3.1 Asset-Allocation Framework and Predictive Regression

One approach to assessing the economic significance of return predictability using $\hat{cay}$ is to document the horizon effects (if any) in optimal asset-allocation decisions made by a utility-maximizing investor. The investor chooses weightings in a portfolio comprising one risky asset (a stock portfolio) and one riskless asset, then holds this portfolio for K periods. Portfolio weights are chosen to maximize the expected power utility of terminal wealth. Since the optimal allocation to risky stocks is horizon dependent if stock returns are predictable, horizon effects can be viewed as a measure of the economic significance of return predictability.

The key input to 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. In this paper, we follow Kandel and Stambaugh (1996) and Barberis (2000) who use a vector autoregression (VAR) to model the joint dynamics of returns and the predictor variable(s). Consider a predictive regression with $\hat{cay}$ as the sole predictor. The system of equations estimated is:

$$\begin{aligned} r_t &= b_{10} + b_{11}\hat{cay}_{t-1} + e_{1t} \\ \hat{cay}_t &= b_{20} + b_{21}\hat{cay}_{t-1} + e_{2t}. \end{aligned} \tag{6}$$

The equations in (6) can be viewed as a VAR with the coefficient on lagged returns r_{t-1} restricted to zero in the first equation. For purposes of the current discussion, the variance-covariance matrix of the innovations is Σ, with elements σ_{ij}.

The results of Kandel and Stambaugh (1996) and Barberis (2000) demonstrate that naive inferences about the economic significance of predictability based on simple regression R^2 may be misleading, as it is important to account for the correlation between innovations in returns and the predictor(s). Barberis provides a particularly striking example of how very modest predictability over short horizons may dramatically alter the risk-return tradeoff, and consequently, the asset allocation of risk averse investors over longer horizons.

In particular, stocks are more appealing over longer horizons if the conditional variance of stock returns grow less than linearly with horizon. Barberis uses a simple two-period example to illustrate the mathematics and explain the intuition in terms of negative correlation between innovations in returns and the (dividend yield) predictor. That is, a negative shock to dividend yields is likely to be associated with a positive shock to expected return if $\sigma_{12} < 0$. A lower dividend yield implies a lower expected return. This sequence of a high (low) realized return followed by a low (high) expected return implies that realized returns will exhibit a degree of negative serial correlation, and the conditional variance of cumulative returns does not scale with horizon.[7]

3.2 Parameters and Uncertainty

Having estimated both the cointegrating regression and predictive VAR, the Bayesian posterior densities capture relevant information relating to the parameters of each model.

This information allows the analysis of return predictability to proceed in one of two ways. First, point estimates of θ_0 can be used to construct estimates of $\hat{cay}$. Conditioning on $\hat{cay}$, the estimation risk associated with the parameters of the predictive vector autoregression θ_v can be accounted for in optimal asset-allocation decisions.

Specifically:

$$p(r_{t+j}|r_t, c\hat{a}y_{t-1}) = \int p(r_{t+j}|r_t, \theta_v, c\hat{a}y_{t-1})p(\theta_v|\mathcal{D}, c\hat{a}y_{t-1})d\theta_v,$$

where we condition on $c\hat{a}y_{t-1}$, or equivalently, c, a, y and the cointegrating parameters θ_0.

An alternative approach is to incorporate the uncertainty surrounding both θ_0 and θ_v. In doing so, the optimal asset-allocation decision is conditioned only on the data $\mathcal{D} \equiv (c, a, y, r)$ and model specification.

Using the numerical posterior estimates based on the approach described in section I, we consider in section 5.2.2 return predictability conditional on the data only. That is,

$$p(r_{t+j}|r_t, c, a, y) \quad = \quad \int\int p(r_{t+j}|\mathcal{D}, \theta_v, \theta_0)p(\theta_v, \theta_0|\mathcal{D})d\theta_v d\theta_0.$$

4 Data

4.1 Consumption, Asset Wealth and Income

The data employed in this study are drawn from a number of sources. The construction of variables and macroeconomic time-series data for c, a and y are identical to LL, albeit over an extended time period extending from Q1 1951 to Q1 2003.[8]

4.2 Equity Return Data

In empirical applications we investigate the predictive properties of *cay* using quarterly returns on the S&P500 index. We use the quarterly dividend yield on the S&P500 as a predictive benchmark.

To approximate the risk-free rate we calculate the quarterly return on an investment in 1-month U.S. Treasury bills.

All returns are continuously compounded.

4.3 The Index of Consumer Sentiment

The Index of Consumer Sentiment (ICS), devised in the 1940's by George Katona and published since 1960 by the University of Michigan, provides a composite measure of consumers' perceptions of their 'ability and willingness to buy'. Specifically, the index seeks to capture consumers' ability and willingness to undertake discretionary expenditures.[9]

The ICS is constructed based on responses to the following questions:

1. *We are interested in how people are getting along financially these days. Would you say that you feel that you (and your family living there) are better off or worse off financially than you were a year ago? Why do you say so?*

2. *Now looking ahead - do you think that a year from now you (and your family living there) will be better off financially, or worse off, or just about the same as now?*

3. *Now turning to business conditions in the country as a whole - do you think that during the next 12 months we'll have good times financially, or bad times, or what?*

4. *Looking ahead, which would you say is more likely - that in the country as a whole we'll have continuous good times during the next five years or so, or that we'll have periods of widespread unemployment or depression, or what?*

5. *About the big things people buy for their homes - such as a refrigerator, a stove, television and things like that. Generally speaking, do you think that now is a good or a bad time for people to buy major household items? Why do you say so?*

The index is calculated as follows:

$$ICS_t = \sum_{j=1}^{5}[P_{jt}^f - P_{jt}^u]100 + 100,$$

where P_{jt}^f is the proportion of the sample giving favourable responses to question j at time t and P_{jt}^u is the proportion of the sample giving unfavorable responses to question j at time t.

The University of Michigan surveys a representative sample of at least 500 US households on a monthly basis by telephone. An independent sample of households is drawn each month, hence, the quarterly surveys are based on a minimum survey sample size of 1500.

Time series of the sentiment index and its components are made freely available on the web by the University of Michigan.

5 Results

Figure 1 presents a standardized plot of the ICS and $\hat{cay}$ - based on point estimates of the cointegrating parameters. The most striking feature of the data is the strong negative (-27%) contemporaneous correlation between the measured sentiment and consumer behavior as summarized by trend deviation, $\hat{cay}_t$. Stronger still is the negative (-67%) contemporaneous correlation between the dividend yield[10] and consumer sentiment evident in figure 2.

Negative correlation between consumer confidence and the dividend yield (or $\hat{cay}_t$) is consistent with time varying risk aversion as per Campbell and Cochrane (1999). Consumption rises above habit and risk aversion declines during consumption booms. The decline in risk aversion leads to a greater demand for risky assets and a decline in expected returns (risk premia). Hence, booms are times of rising consumption, but declining ratios of consumption to wealth.

5.1 Predictive Regressions

Table 1 reports predictive regressions estimated over the interval from the first quarter of 1960 to the first quarter of 2003. As noted earlier, $\hat{cay}$ is a highly significant predictor of returns with a raw regression R^2 of 8.5% and DY is a marginally significant predictor of return with a regression R^2 of 1.78%. As a stand-alone predictor of the quarterly equity return premium, the ICS appears useless - a regression R^2 that's indistinguishable from zero, and there is no value in adding the ICS to a predictive regression containing $\hat{cay}$ or DY.

Whilst the ICS does not seem to have any stand-alone predictive value, the second last regression in table 1 reveals a significant interaction with $\hat{cay}$. The ratio of consumption to wealth scaled by measured consumer sentiment is statistically significant when added to $\hat{cay}$ - and the adjusted R^2 of the predictive regression jumps from 7.87% to 10.8%. In economic terms, point estimates of the (standardized) raw regression coefficients imply that the predictive implications of $\hat{cay}$ are heavily influenced by the level of measured sentiment.

Ignoring consumer sentiment, based on the shorter estimation interval for the predictive regression, a standard deviation shift in $\hat{cay}$ implies a 261 basis point change in the expected quarterly return premium. If consumer sentiment is neutral, the same standard deviation shift in $\hat{cay}$ implies a 329 basis point change in the expected quarterly return premium. If the ICS is a standard deviation above its historical mean, the standard deviation change in $\hat{cay}$ implies a 169 basis point change in the quarterly return premium. Clearly, the regression results imply that the level of $\hat{cay}$ should be interpreted in light of the ICS during the same period.

Based on the regression evidence in table 1, we observe no corresponding interaction between sentiment and the dividend yield.[11]

5.2 Vector Autoregressions and Parameter Uncertainty

An alternative approach to studying the predictive properties of $\hat{cay}$ and DY is to model the joint dynamics of returns and the predictor(s), assuming they are well described by a VAR. To consider the impact of parameter uncertainty we adopt a Bayesian perspective along the lines of Kandel and Stambaugh (1996) and Barberis (2000).

Table 2 reports the posterior mean of the predictive VAR parameters for quarterly returns using $\hat{cay}$ and the lagged quarterly dividend yield as predictors.[12] The most significant points of contrast between the quarterly VAR results in table 2 and the monthly VAR results reported by Barberis are the correlations between innovations in returns and the predictors. At -58% and -69% respectively, the mean correlations between innovations in quarterly returns and $\hat{cay}$, and innovations in quarterly returns and the dividend yield are substantially smaller in magnitude than the -94% for monthly returns and the dividend yield reported in Barberis (2000) over a somewhat shorter sample period.[13]

5.2.1 Predictive VAR Parameters Conditional on *cay*

The economic implications of the quarterly parameter values can be illustrated in terms of the buy and hold asset allocation problem Barberis (2000) uses to demonstrate horizon effects. Figures 4 - 5 graph the optimal allocations employing both predictive variables using two different levels of risk aversion for a constant relative risk aversion investor.[14]

The solid line with pentagrams in panel A of figure 4 shows that a power utility investor with risk aversion of $A = 5$ and an investment horizon of one quarter allocates just over 60% of wealth to equity (and the remainder to the short term riskless asset) when using the lagged quarterly dividend yield to predict returns. The same investor with a 10-year horizon allocates over 72% of wealth to equity, conditional on the posterior mean of the VAR parameters. The importance of this horizon effect weakens significantly if risk aversion is higher. For example, in panel B of figure 4 reveals that an investor with $A = 20$ invests allocates 15% of wealth to equity over a 1-quarter horizon, and approximately 17% over a 10-year horizon.

The corresponding results for an investor using $\hat{cay}$ as the predictor can be seen in figure 5. The solid line with pentagrams in panel A of figure 5 shows that an investor with $A = 5$ invests approximately 65% of wealth in equity for a 1-quarter horizon, and 76% for a 10-year horizon. Whilst (again) the horizon effect is small relative to Barberis's, the notable aspect of the result is the initial climb and subsequent decline in the allocation to equity. When $\hat{cay}$ is used in the predictive VAR, the optimal allocation of a buy and hold investor does not increase monotonically with horizon (conditional on the posterior mean of the VAR parameters). The allocation to equity peaks at 85% (for 10- and 11-quarter horizons) and then declines to 76% as horizon increases to 10 years. A similar pattern (on a smaller scale) is true when $A = 20$.

The pattern of asset allocation conditional on $\hat{cay}$ can be explained in terms of the implied variation in annualized Sharpe ratio by investment horizon. Consistent with the pattern in annualized volatility in figure 3 and equity allocation, the Sharpe ratio initially rises and then falls as investment horizon grows.

The dotted lines with pentagrams in figures 4 - 5 plot optimal allocations conditioned on time t (initial) values of each predictive variable a standard deviation above their historical mean. In terms of incremental impact on asset allocation, the initial value of each predictor has similar properties. For example, in panel B of figure 4, it can be seen that when $A = 20$ the initial allocation to equity increases by almost two thirds when the dividend yield is a standard deviation above its historical average relative to the benchmark historical average case (solid line with pentagrams).

As can be seen in panel B figure 5, when $A = 20$ and $c\hat{a}y$ is one standard deviation above its historical mean, the initial impact on equity allocation is even more dramatic. Relative to the historical mean (benchmark) case, the allocation to equity more than doubles. However, the time variation in periodic volatility implies that the long term allocation to equity falls below the initial 1-quarter allocation. That is, an investor with a 1-quarter horizon allocates 36% of wealth to equity, but less than 23% over a 10-year horizon - again, after an initial increase in weighting with horizon, and a subsequent decline.

By contrast, when $c\hat{a}y$ is set to a value one standard deviation below its mean, the optimal allocation for the same $A = 20$ investor increases with horizon from a value of zero for a 1-quarter horizon, to over 12% for a 10-year horizon (dashed line with pentagrams in figure 5). Negative standard deviation shocks to the inital value of the dividend yield have significant but less dramatic effects on the initial allocation to equity relative to the benchmark case. Further, the horizon effects are still present but very weak.

A common theme to the results in figures 4 - 5 is that estimation risk lowers the allocation to equity - regardless of predictor or level of risk aversion. Horizon effects associated with the dividend yield are eliminated. In fact, the plots without pentagrams in both panels of figure 4 suggest that estimation risk dominates, and all else being equal, the optimal static allocation to equity declines with investment horizon. However, it is also evident, that even after allowing for estimation error in VAR parameters, a standard deviation shock to the dividend yield retains an economically substantial effect on optimal allocation to equity.

Whilst the horizon effects associated with $c\hat{a}y$ are weakened when we allow for estimation risk in VAR parameters, optimal equity allocations remain sensitive to $c\hat{a}y$ shocks - particularly over horizons of 1 to 20 quarters.

In summary, the combined effects of estimation error and VAR parameter values lead to a similar conclusion for both DY and $c\hat{a}y$ over longer horizons. That is, in this simple stylized setting, long horizon investment weights are quite similar to short horizon weights at a given level of risk aversion, given equivalent DY or $c\hat{a}y$ signals, though for different reasons. The mean predictability associated with DY is swamped by estimation risk, whilst the horizon effects associated with $c\hat{a}y$ are dampened by the combined effects of estimation risk and the correlation structure of VAR innovations.

5.2.2 Accommodating Uncertainty Associated with Cointegrating Parameters

Posterior Distribution of Cointegrating Parameters

In the cointegrating regression (I.6), there are two primary sources of uncertainty affecting the long-run dynamics of the relation between (log) consumption, asset wealth and income: (i) the cointegrating vector β and (ii) the adjustment coefficients α. Figures 6 and 7 provide a graphical summary of the uncertainty in parameters governing the long-run dynamics, and table 3 provides basic summary statistics.[15] As may be expected, the standard deviation of the cointegrating vector components are small relative to the means. Also consistent with expectations, the posterior standard deviation of the adjustment coefficients for consumption and income are large - of the same order of magnitude or larger than their expected values - whilst the dispersion of the adjustment coefficient on asset wealth is relatively small.

LL use dynamic least squares based point estimates of the cointegrating parameters to document departures from the steady state ratios of income to wealth. Their estimates (based on a slightly shorter sample period) suggest that asset wealth accounts for 34% of aggregate wealth. The posterior shown in figure 6 suggests a mean value of 31% with a standard deviation of 5.5%. Before we turn to the question of whether this translates to an economically significant risk in forecasting returns, we can summarize this uncertainty in terms of $\hat{cay}$.

Figure 8 illustrates the posterior mean and 90% highest posterior density (HPD) interval of the $\hat{cay}$ posterior implied by β and α. In terms of the posterior mean, the 90% HPD interval ranges from 0.35 to over 2.95 standardized units. For example, the most recent 90% HPD interval is bounded by standardized $\hat{cay}$ values of 0.096 and 1.58. This is a wide dispersion in terms of the predictability documented in LL - suggesting that a single (time series) standard deviation shift in $\hat{cay}$ corresponds to a 220 basis point change in quarterly real returns on the S&P 500. The histogram in figure 9 summarizes the posterior distribution of cay based on $p(\theta_0|\mathcal{D})$ at the end of the sample (2003, Q1). In the absence of non-sample information, there is substantial uncertainty about cay: its posterior mean 0.65 and its posterior standard deviation is 0.4 (standardized units). We now examine the effects of this uncertainty on cay's predictive properties.

Revisiting the Predictive Regression

To get a sense of how the uncertainty captured in figures 8 and 9 alters the most basic measure of predictability we estimate a predictive regression of the return premium conditional on the $\hat{cay}$ series implied by each draw from $p(\theta_0|\mathcal{D})$:

$$\begin{array}{rcl} r_t & = & 0.018 \; + \; \underset{(0.0031)}{0.0243} \, \hat{cay}_{t-1} \; + \; e_t. \\ & & \qquad\qquad\qquad (R^2 = 0.087, \sigma_{R^2} = 0.018) \end{array} \tag{7}$$

The point estimates reported in equation (7) are the mean regression coefficients obtained from 100,000 regressions.[16] The standard deviation of 0.0031 for the coefficient of $\hat{cay}_{t-1}$ and 0.018 for the R^2 reflects the variability attributable to uncertainty in $\hat{cay}$ parameters. Given the often large dispersion in the posterior values of *cay* at various points in time, the simple regression estimates of predictability are remarkably robust: neither the size of the predictive regression slope coefficient nor the R^2 appear sensitive to the uncertainty in *cay*. The variability of *cay* does not translate to significant variation in the predictive regression results.

We now turn to the question of whether this holds true for longer horizon inference based on a predictive vector autoregression.

Implied Long Horizon R^2 Conditional on DY and $\hat{cay}$

The long horizon R^2 suggested by Hodrick (1992) summarizes the impact of short horizon predictability based on a VAR. In the current application, the long horizon R^2 is implied by a first order VAR. LL report the implied long horizon R^2 statistics based on point estimates of the VAR parameters, conditional on the cointegration parameters using multiple predictors. Here, we focus on the long horizon predictability associated with *cay* and the impact of uncertainty associated with the VAR and cointegrating parameters. We benchmark our findings against the corresponding metric computed with respect to DY.

Figures 10 and 11 graph the 90% highest posterior density (HPD) intervals of the long horizon R^2s associated with predictive VARs employing DY and *cay* respectively. Since Hodrick's R^2 is a function of the VAR parameters, we obtain its posterior using the draws from $p(\theta_v|\theta_0, \mathcal{D})$ and $p(\theta_v|\mathcal{D})$ as a measure of the incremental impact of uncertainty associated with the cointegrating parameters.

The uncertainty associated with the predictability implied by a VAR based on DY, as measured by the 90% HPD of the implied R^2, is large. As can be seen in figure 10, at a two year horizon the spread of the implied R^2 exceeds 30%, at a five year horizon the corresponding dispersion exceeds 50% and at a horizon of ten years the range of the HPD exceeds 70%.

The shaded area in figure 11 corresponds to the region covered by the 90% HPD of the long horizon R^2 implied by the VAR based on $\hat{cay}$. The dispersion of the long horizon R^2 based on $\hat{cay}$ is much smaller than the dispersion associated with DY. The widest HPD is less than 30%, and it occurs at a horizon of two years. After two years, the posterior median of the implied R^2 and the associated uncertainty decline. At a horizon of ten years, the posterior median is 10% and the 90% HPD has a range of approximately 15%.

Implied Long Horizon R^2 Conditional on Data Only

Instead of conditioning predictions on point estimates of the cointegrating parameters, we utilize the sampling scheme described in appendix I and extend the simple VAR analysis to account for the fact that $\hat{cay}$ is itself estimated. The shaded area in figure 11 is based on $p(\theta_v|\mathcal{D}, \theta_0)$, but we now require draws from $p(\theta_v, \theta_0|\mathcal{D})$.

To draw from $p(\theta_v, \theta_0|\mathcal{D})$ we sample as follows.

1. Draw from:

   $$p(B|\Sigma, \theta_0, \mathcal{D}, c_{t,pred}, a_{t,pred}, y_{t,pred})$$

   and

   $$p(\Sigma|B, \theta_0, \mathcal{D}, c_{t,pred}, a_{t,pred}, y_{t,pred})$$

   as per the sampling scheme in Barberis (2000). This yields a draw from $p(B, \Sigma|\theta_0, \mathcal{D}, c_{t,pred}, a_{t,pred}, y_{t,pred})$. The first draw is conditioned on an arbitrary θ_0, and each subsequent draw $[i+1]$ is conditioned $\theta_0^{[i]}$.

2. Draw from $p(\theta_0|B, \Sigma, \mathcal{D}, c_{t,pred}, a_{t,pred}, y_{t,pred}) = p(\theta_0|\mathcal{D})$ as described in appendix I conditioning on $B^{[i+1]}$ and $\Sigma^{[i+1]}$.[17]

The dashed lines in figure 11 trace out the 90% HPD of the implied R^2 upon accounting for the uncertainty associated with both the cointegrating parameters and the VAR parameters. In terms of the HPD, the incremental effect of this uncertainty widens the HPD by 2-8%, increasing with horizon. However, even after accounting for the uncertainty in both the VAR and cointegrating parameters, the uncertainty associated with the predictive properties of *cay* is much smaller than the corresponding uncertainty associated with predictions based on DY.

Whilst the evidence of in-sample return predictability based on *cay* appears robust to the uncertainty in the cointegration parameters, the dispersion of the *cay* posterior at a given point in time has implications for the usefulness of *cay* as a predictive signal. Returning to the stylized asset allocation problem, the asset allocation decision is no longer conditioned on $\widehat{cay}$, as the investor must now account for the noise in cointegration parameters and the attendant uncertainty about the departure from equilibrium. As shown in figures 8 and 9, given values of c, a, and y are consistent with a wide range of residuals associated with plausible cointegrating parameter draws.

For the sake of our example we condition our estimates of the predictive distributions on historical values of c, a and y that correspond most closely to a standard deviation shift in $\widehat{cay}$, conditional on the posterior mean of the cointegration parameters. Specifically, treating the posterior mean of the cointegration parameters as fixed, an investor in the third quarter of 1953 would have been confronted with a value of $\widehat{cay}$ approximately one standardized unit above its mean. In the third quarter of 2002 the same investor would have been confronted with $\widehat{cay}$ one standardized unit below its mean. Accordingly, in figure 12 we present asset allocations conditional on observations of c, a and y as at Q3 1953 and Q3 2002 to analyze the case where an investor relying on point estimates of the cointegration parameters would consider $\widehat{cay}$ to be "high" and "low" respectively.

In strong contrast to the results in figure 5, equity allocations now appear relatively insensitive to signals in the data on consumption, asset wealth and income. When the coefficient of risk aversion is low, the difference between allocation to equity given a high signal (solid line) and a low signal (solid line with pentagrams) is observable, but only a small fraction of what we observed in panel A of figure 5. When the coefficient of risk aversion is relatively high, asset allocation is basically independent of the signal. These result suggest that even in light of apparently robust historical evidence of short horizon predictability based on $\widehat{cay}$, data on c, a and y do not necessarily contain economically significant signals once we account for the combined effects of estimation risk in modeling the predictive and cointegrating relations.

5.2.3 The Impact of Sentiment

The regression estimates in table 1 suggest a statistically and economically significant interaction between sentiment and $\widehat{cay}$. To better understand the importance of this result, we consider the joint dynamics of the return premium, $\widehat{cay}$ and the interaction between $\widehat{cay}$ and sentiment modeled as a first order VAR.

Figure 13 illustrates the impact of sentiment in the context of the investment decision considered earlier.

A shock to $\widehat{cay}$ has very different implications for short to medium term asset allocation when we allow for effect of sentiment. Defining optimism (pessimism) as a positive (negative) standard deviation from mean sentiment we can see that a pessimistic investor's short term allocation to equity is almost double that of an optimistic investor's in response to a positive standard deviation shock to $\hat{cay}$. Negative shocks to $\hat{cay}$ have a slightly larger effect on the asset allocation of a pessimistic investor than the corresponding optimistic investor. These results suggest that the findings of significant interaction between the ICS and $\hat{cay}$ are robust. Shifts in $\hat{cay}$ should be considered in light of sentiment, as measured by the ICS.

These findings are conditional on a point estimate of $\hat{cay}$ and must of course be interpreted in light of the results in section 5.2.2.[18]

5.3 A Note on Sentiment, *cay* and Causality

Preliminary investigations of Granger causality using the components of the ICS reveal evidence of bi-directional causality.

Consumers beliefs about current buying conditions (as measured by the index based on responses to question 5) appear to Granger cause $\hat{cay}$. That is we reject the null of no causality in a predictive VAR using all ICS components across all specifications of lag length considered (models with 1-6 predictor lags were estimated). There is some evidence of $\hat{cay}$ Granger causing business conditions - however this result is sensitive to the specification of VAR lag length. The same is true of 12 month business conditions (based on responses to question 3) - the result is sensitive to VAR lag length specification.

Overall, we reject the null of no causality when we use $\hat{cay}$ to forecast the ICS at all specifications of lag length considered. There is no evidence of the ICS Granger causing $\hat{cay}$ at any conventional level of statistical significance at any specification of lag length.[19]

6 Conclusions

We've studied the implications of consumer expectations as measured by LL's $\hat{cay}$ and the University of Michigan's ICS for expected equity returns.

Consistent with LL, we find strong evidence of return predictability based on departures from an intertemporal budget constraint, as measured by $\hat{cay}$. The finding is robust in the sense that the predictability is evident after accounting for various sources of estimation risk. However, the apparent predictability does not necessarily give rise to useful signals based on shifts in aggregate consumption, income and asset wealth due to estimation risk associated with $\hat{cay}$ at a given point in time.

We find that the survey based measure of expectations complements the behavioral measure, but has no apparent stand-alone predictive value in forecasting equity returns.

References

Barberis, N., 2000, "Investing for the Long Run when Returns are Predictable," *Journal of Finance*, 55, 225–264.

Bauwens, L., and P. Giot, 1998, "A Gibbs Sampler Approach to Cointegration," *Computational Statistics*, 13, 339–368.

Bauwens, L., and M. Lubrano, 1996, "Identification Restrictions and Posterior Densities in Cointegrated Gaussian VAR Systems," in *Advances in Econometrics*. JAI Press, Greenwich, CT, vol. 11, pp. 3–28.

Brennan, M., and Y. Xia, 2002, "tay's as good as cay," Working Paper.

Campbell, J., and J. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behaviour," *Journal of Political Economy*, 107, 205–251.

Cochrane, J., 1999, "Portfolio Advice for a Multifactor World," in *Economic Perspectives*. Federal Reserve Bank of Chicago.

Curtin, R., 2000, "Psychology and Macroeconomics: Fifty Years of the Surverys of Consumers," working paper, University of Michigan.

——— , 2002, "Surveys of Consumers: Theory, Methods and Interpretation," working paper, University of Michigan.

Geweke, J., 1996, "Bayesian Reduced Rank Regression in Econometrics," *Journal of Econometrics*, 75, 121–146.

Hahn, J., and H. Lee, 2002, "On the Estimation of the Consumption-Wealth Ratio: Cointegrating Parameter Instability and its Implications for Stock Return Forecasting," Working Paper, GSB Columbia University.

Hodrick, R., 1992, "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 5, 357–386.

Johansen, S., 1988, "Statistical Analysis of Cointegrating Vectors," *Journal of Economic Dynamics and Control*, 12, 231–254.

——— , 1991, "Estimation and Hypothesis Testing of Cointegrating Vectors in Gaussian Vector Autoregressive Models," *Econometrica*, 59, 1551–1580.

Kandel, S., and R. Stambaugh, 1996, "On the Predictability of Stock Returns: An Asset Allocation Perspective," *Journal of Finance*, 51, 385–424.

Kleibergen, F. R., and R. Paap, 2002, "Priors, Posteriors and Bayes Factors for a Bayesian Analysis of Cointegration," *Journal of Econometrics*, 111, 223–249.

Kleibergen, F. R., and H. K. van Dijk, 1998, "Bayesian Simulataneous Equations Analysis using Reduced Rank Structures," *Econometric Theory*, 14, 701–743.

Lettau, M., and S. Ludvigson, 2001, "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56, 815–849.

——— , 2002, "tay's as good as cay: a reply," Working Paper.

Luukkonen, R., A. Ripatti, and P. Saikkonen, 1999, "Valid Bayesian Estimation of the Cointegrating Regression Model," *Journal of Business and Economic Statistics*, 17, 195–204.

Matsusaka, J., and A. Sbordone, 1995, "Consumer Confidence and Economic Fluctuations," *Economic Inquiry*, XXXIII, 296–318.

Rudd, J., and K. Whelan, 2002, "A Note on the Cointegration of Consumption, Income and Wealth," Working Paper.

Stock, J., and M. Watson, 1993, "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 6, 783–820.

Strachan, R., 2003, "Valid Bayesian Estimation of the Cointegrating Regression Model," *Journal of Business and Economic Statistics*, 23(1), 185–195.

Villani, M., 2001a, "Bayesian Prediction with a Cointegrated Vector Autoregression," *International Journal of Forecasting*, pp. 585–605.

——— , 2001b, "Bayesian Reference Analysis of Cointegration," Working Paper, University of Stockholm.

Zellner, A., 1971, *An Introduction to Bayesian Inference in Econometrics*. John Wiley & Sons Inc, New York.

Notes

[1]See for example Cochrane (1999).

[2]This is true even though we use the full span of he sample to estimate them.

[3]The estimated predictive regressions is:

$$\begin{array}{rcl} r_t & = & \underset{(3.10)}{0.018} + \underset{(1.96)}{0.0282}\, DY_{t-1} + e_t \qquad (R^2 = 0.0186) \end{array}$$

[4]Lettau and Ludvigson (2002), page 4: 'We conclude that, while the explanation for *why* $c\hat{a}y$ forecasts stock returns is open to interpretation, the findings of predictability per-se are not.'

[5]Refer to Curtin (2000) for further discussion this dichotomous view of rationality in the social sciences and references to further readings.

[6]Refer to Curtin (2002) for detailed discussion of interpretation and performance of the expectation measures. It is interesting to note that the discretionary expenditures captured by the consumer sentiment surveys are excluded from the consumption figures used to construct $c\hat{a}y_{t-1}$ on the grounds that expenditures on durable goods are replacements and additions to a stock rather than a 'service flow' from the existing stock (LL, page 822).

[7]The solid line plot in figure 3 illustrates the effect using the posterior mean of the parameter estimates reported by Barberis (2000) (reproduced in Table 2 for comparative purposes) employing the monthly dividend yield as a predictive variable over the 1952-1995 interval. The conditional periodic volatility of return declines monotonically with investment horizons ranging from a quarter to ten years. If expected returns are positive and grow linearly with horizon, the implied Sharpe ratio of equity increases with horizon and becomes more attractive to risk averse investors.

[8]The data is available on Martin Lettau's web page:

http://pages.stern.nyu.edu/~mlettau/index.htm.

These data are quarterly, seasonally adjusted, per capita variables measured in 1992 chain weighted dollars. The reader is referred to the appendix in Lettau and Ludvigson (2001) for details on variable construction and data sources.

[9]Refer to Curtin (2002) for detailed discussion of interpretation and performance of the expectation measures. It is interesting to note that the discretionary expenditures

captured by the consumer sentiment surveys are excluded from the consumption figures used to construct $\hat{cay}$ on the grounds that expenditures on durable goods are replacements and additions to a stock rather than a 'service flow' from the existing stock (LL, page 822).

[10]The dividend yield is the quarterly mean dividend yield on the S&P500.

[11]Estimates of the predictive regressions in table 1 using real returns on the S&P 500 yield almost identical results.

[12]The sampling scheme to estimate the predictive VAR is described in Barberis (2000). The conditional posteriors are based on diffuse priors.

[13]This is important in in the context of the asset allocation problem considered in this paper, as the result on page 245 of Barberis (2000) is not as general as he suggests. The dashed graph on 3 plots the periodic volatility when the covariance between innovations is halved, and all other parameters are fixed at their original values. Whilst periodic volatility initially falls with horizon, it flattens out, turns, and begins to increase with horizon when the correlation between innovations in returns and the predictor is lowered.

Additionally, note that the quarterly dividend yield exhibits a very high level of persistence - similar to the monthly result and substantially higher than that of $\hat{cay}$.

[14]The investor is assumed to have power utility over terminal wealth W at each horizon: $U(W) = \frac{W^{1-A}}{1-A}$. To facilitate comparison of the current results with those reported in Barberis (2000) we consider: $A = 5$ & 20.

[15]The reported results do not seem sensitive to the specification of k.

[16]One regression for each draw of the cointegrating parameter set.

[17]To sample from $p(r_{t+1}|\mathcal{D}, c_{t,pred}, a_{t,pred}, y_{t,pred})$, an additional step is added: Draw from $p(r_{t+1}|B, \Sigma, \theta_0, \mathcal{D}, c_{t,pred}, a_{t,pred}, y_{t,pred})$ conditioning on $B^{[i+1]}$, $\Sigma^{[i+1]}$ and $\theta_0^{[i+1]}$. The $r_{t+1}^{[i]}$ from this sampler are draws from the marginal (predictive) posterior $p(r_{t+1}|\mathcal{D}, B, \Sigma, \theta_0, c_{t,pred}, a_{t,pred}, y_{t,pred})$.

[18]We have estimated the analog of (7) and confirmed that the finding of predictability is not particular to the point estimate $\hat{cay}$.

[19]These results are available on request. Including the quarterly S&P500 return premium in the predictive VAR is interesting. The return premium Granger causes everything - $\hat{cay}$ and the ICS. Once we account for the market premium, there is no evidence of $\hat{cay}$ Granger causing ICS at conventional levels of significance.

I Estimating the Cointegrating Regression

There are several recent Bayesian studies of cointegration, including Bauwens and Giot (1998), Bauwens and Lubrano (1996), Geweke (1996), Kleibergen and Paap (2002), Kleibergen and van Dijk (1998), Luukkonen, Ripatti, and Saikkonen (1999), Strachan (2003), Villani (2001b) and Villani (2001a). In this paper, our analysis of cointegration between c, a and y adopts the estimation approach presented in Strachan (2003) which ensures valid Bayesian estimation of cointegrating parameters.

The Strachan (2003) approach can be thought of as the Bayesian analogue of the Johansen (1988) and Johansen (1991) approach to identification in a classical setting. Strachan provides a particularly clear exposition of why certain identifying restrictions on the cointegrating vector may not be innocuous. Linear identifying restrictions restrict the length of one coefficient such that the directions of the cointegration space are restricted. Strachan's estimation procedure achieves identification by restricting the length of the cointegrating vector without placing arbitrary restrictions on the angle of the space spanned by it.[20] The framework also provides the foundations of finite sample inference on the validity of identifying restrictions and the rank of the long-run impact matrix.

Consider the p-dimensional vector-valued process $x_{t=1\ldots T}$:

$$x_t = \sum_{i=1}^{K} \Pi_i x_{t-i} + \Phi d_t + \varepsilon_t. \tag{I.1}$$

Writing (I.1) as an Error Correction model (ECM), we get:

$$\Delta x_t = \Pi x_{t-1} + \sum_{i=1}^{k-1} \Gamma_i \Delta x_{t-i} + \Phi d_t + \varepsilon_t, \tag{I.2}$$

where $\Delta x_{t-i} = x_{t-i} - x_{t-i-1}$, $\Pi = \sum_{i=1}^{K} \Pi_i - I_p$ is a $p \times p$ long-run impact matrix, and $\Gamma_i = -\sum_{j=1+1}^{K} \Pi_j$ for $i = 1, \ldots, k-1$ is a $p \times p$ matrix of coefficients governing the short-run dynamics, d_t is a $w \times 1$ vector of deterministic variables, Φ is the $p \times w$ coefficient matrix and ε_t is a zero-mean vector of errors with covariance Ω.

If all p time series comprising x_t are $I(1)$ (integrated of order 1) and $\text{rank}(\Pi) = r < p$, then r linear combinations of the series are stationary and the time-series are said to be cointegrated. Further, if $\text{rank}(\Pi) = r$, then Π can be written as the product of two $p \times r$ full rank matrices α and β:

$$\Delta x_t = \alpha\beta' x_{t-1} + \sum_{i=1}^{k-1} \Gamma_i \Delta x_{t-i} + \Phi d_t + \varepsilon_t. \tag{I.3}$$

We can interpret $\beta' x_{t-1} - E(\beta' x_{t-1})$ as deviations from the r equilibria and α as a matrix of coefficients governing the adjustment to equilibrium after a disturbance. Each column of the matrix β is a cointegrating vector and the matrix is determined up to an arbitrary linear combination of its columns. To overcome the indeterminacy, the following normalization is commonly applied:

$$\beta = \begin{bmatrix} I_r \\ \beta^* \end{bmatrix}, \tag{I.4}$$

such that β^* is $(p-r)\times r$ and fully identified. The normalization can only be applied if each of the r first components of x_t enter at least one of the cointegrating relations.

Strachan (2003) achieves identification by nesting a reduced rank model within a general full rank model with a well-behaved posterior distribution. The full rank model is transformed and the restriction is paramaterized conditional on the rank r. Estimation involves taking a singular value decomposition of Π with respect to its estimated variance-covariance matrix:

$$\Pi = \beta\alpha + S_{11}^{-1}\beta_\perp \lambda \alpha_\perp \tilde{\Sigma}, \tag{I.5}$$

where $\tilde{\Sigma} = S_{00} - S_{01}S_{11}^{-1}S_{10}$, S_{ij} are moment matrices when a diffuse prior is used, $\beta_\perp$ and $\alpha'_\perp$ are $p \times (p-r)$ matrices orthogonal to β and α' respectively.[21]

By re-writing (I.2) in terms of the singular value decomposition:

$$\Delta x_t = \beta\alpha x_{t-1} + S_{11}^{-1}\beta_\perp \lambda \alpha_\perp \tilde{\Sigma} x_{t-1} + \sum_{i=1}^{k-1} \Gamma_i \Delta x_{t-i} + \Phi d_t + \varepsilon_t, \tag{I.6}$$

it can be seen that cointegration occurs when $\lambda = 0$. Letting $\mathcal{D}$ denote the sample data, we can write our focus of interest, the (joint) posterior of model parameters $p(\theta|\mathcal{D})$, conditional on $r < p$ with $\lambda = 0$ in general terms:

$$p(\theta_0|\mathcal{D}) = p(\theta|\mathcal{D})|_{\lambda=0}. \tag{I.7}$$

With a view to 'letting the data speak', we choose diffuse priors for all model parameters. As a result, a Metropolis-Hastings (M-H) sampling scheme is required to estimate (I.7) since the conditional posteriors of α and β are not standard (known) densities. The M-H algorithm generates samples from the reduced rank model using a candidate generating density based on the approximating full rank model.

We impose the linear restriction (I.4) within the identifying restrictions of Strachan (2003) and implement the M-H sampling scheme using the approach detailed in Kleibergen and Paap (2002) as follows:[22]

1. Draw Ω^{-1}. As in the standard multivariate regression model (see, for example, Zellner (1971)), a non-informative conjugate prior implies $p(\Omega^{-1}|\alpha,\beta,\Gamma^*,r,\mathcal{D}) = \mathcal{W}[(\varepsilon'\varepsilon + A)^{-1}, t+q]$, where $\varepsilon = (\varepsilon_1 \quad \varepsilon_2 \ \dots \ \varepsilon_T)'$, $\Gamma^* = (\Gamma \ \Phi)'$, and $\mathcal{W}(A,q)$ is the Wishart prior. To keep the prior non-informative, we set $q = 3$, and $A = 0_p$.

2. Draw Γ^*. If $p(\Gamma|\alpha,\beta,\Omega,r) \propto C$, we know from the standard multivariate regression model that:

$$p(\Gamma|\alpha,\beta,\Omega,r,\mathcal{D}) = \mathcal{N}\left[vec((Z'Z)^{-1}Z'(\Delta X_t - \beta\alpha X_{t-1})), \Omega \otimes (Z'Z)^{-1}\right],$$

   where $Z = [\Delta X_{t-1} \ \dots \ \Delta X_{t-k} \ D_t]'$, $\Delta X_t = (\Delta x_1 \ \dots \ \Delta x_t)'$, $X_{t-1} = (x_0 \ x_1 \ \dots \ x_{t-1})'$, and $D_t = (d_1 \ d_2 \ \dots \ d_t)'$.

3. Draw Π, or equivalently, α, β and λ. Since $\lambda = 0$ in the case of a reduced rank error correction model, it may seem odd that it arises in the sampling scheme. This occurs because the candidate generating function used in the following M-H step is based on the full rank (or unrestricted) error correction model.

   (a) Assuming a constant prior, a candidate $\Pi^{[i+1]}$ is drawn from $\mathcal{N}\left[vec((X_t'X_t)^{-1}X_t'(\Delta X_t - Z\Gamma^*)), \Omega \otimes (X_t'X_t)^{-1}\right]$. The index i references the draw number; that is, we are at the $[i+1]$-th iteration of the sampler.

   (b) Compute $\alpha^{[i+1]}$, $\beta^{[i+1]}$ and $\lambda^{[i+1]}$ based on the following singular value decomposition: $\Pi^{[i+1]} = U^{[i+1]}S^{[i+1]}V^{[i+1]}$. Refer to pp 188-189 of Strachan (Strachan 2003) for calculation details.

(c) Compute the acceptance probability for draw $[i + 1]$ using the weight calculations detailed on pp 234-235 of Kleibergen and Paap (Kleibergen and Paap 2002). The rejection probability is based on the ratio of the augmented posterior of the restricted error correction model and the posterior of the unrestricted error correction model.[23] If the current draw is rejected, all parameters are set equal to their previously drawn values, that is, $\theta^{[i+1]} = \theta^{[i]}$.

4. Go to step 1 and repeat the loop many times.

This algorithm generates draws from $p(\lambda, \theta_0|\mathcal{D})$, and hence, the draws of $[\alpha,\ \beta,\ \Omega,\ \Gamma]$ are draws from $p(\theta_0|\mathcal{D})$.

Figure 1: Consumer Sentiment and $\hat{cay}$: Q1 1966-Q1 2003
The solid line plots the path of the posterior mean value of *cay*. The dashed line plots the (standardized) Quarterly Consumer Sentiment Index as reported by the University of Michigan.

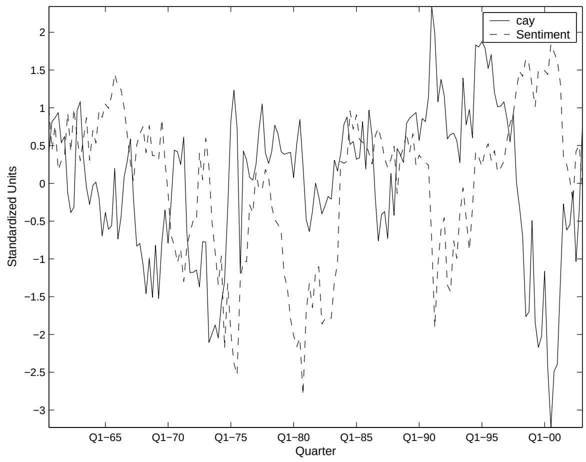

Figure 2: Consumer Sentiment and DY: Q1 1966-Q1 2003
The solid line plots the path of DY. The dashed line plots the (standardized) Quarterly Consumer Sentiment Index as reported by the University of Michigan.

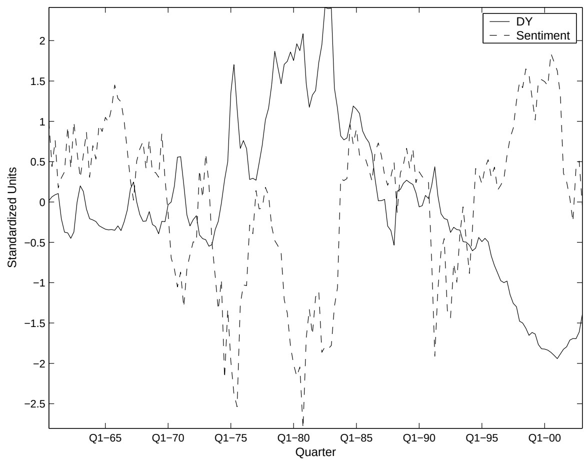

Figure 3: Monthly Return Volatility by Investment Horizon with DY Predictor

The solid line plots the monthly return volatility by investment horizon using the posterior mean of VAR parameters from Barberis (2000) - as reproduced in 2. The correlation between innovations in dividend yield and return is -93.5%. The dashed line plots the same monthly return volatility by investment horizon when the correlation between innovations in returns and dividend yields is -43.5%, all else being equal.

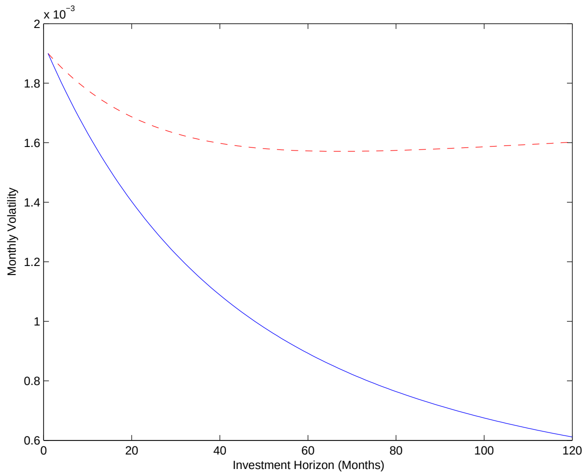

Figure 4: Sensitivity of Buy and Hold Allocation to Initial Dividend Yield

Buy and hold asset allocation by investment horizon: in each chart the solid line with pentagrams plots, by investment horizon, the equity allocation of a power utility investor. In panel A the risk aversion parameter is 5, in panel B the parameter is 20. In each case the investor uses a predictor (quarterly dividend in figure 4, and $\hat{cay}$ in figure 5) to forecast return. The current value of the predictor is set equal to its historical mean and the investor ignores estimation risk in the predictive VAR parameters. The solid line plots the equity allocation of the same investor who accounts for estimation risk in VAR parameters. In figure 5 the values of $\hat{cay}$ are taken as given. The dotted line (dashed) with pentagrams plots the equity allocation by horizon for the case where the current value of the predictor is a standard deviation above (below) its historical mean and estimation risk is ignored. The corresponding plots without pentagrams account for estimation risk in predictive VAR parameters. All results are based on the Q4 1951-Q1 2003 estimation interval and 100,000 draws of the Gibbs sampler. NOTE: scale differs between the panels.

Panel A
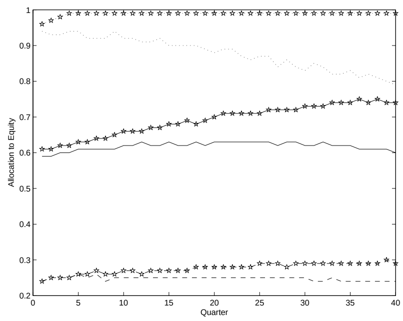

Panel B
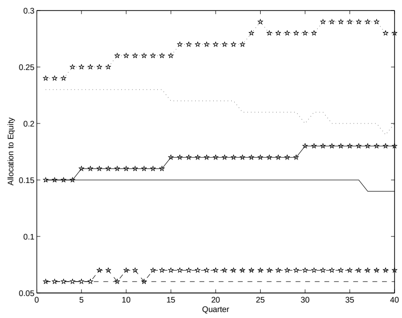

Figure 5: Sensitivity of Buy and Hold Allocation to Initial $\hat{cay}$ Given Wealth Shares

Panel A
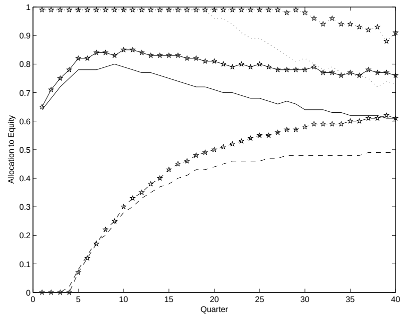

Panel B
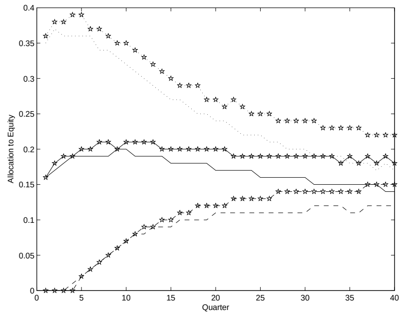

Figure 6: Joint Posterior Distribution of Normalized Cointegrating Coefficients

The posterior $p(\beta_a, \beta_y|c, a, y)$ estimates are based on 100,000 draws of a Metropolis-Hastings sampling scheme with a rejection rate of 20.8%. Parameter settings and priors are described in the caption to table 3.

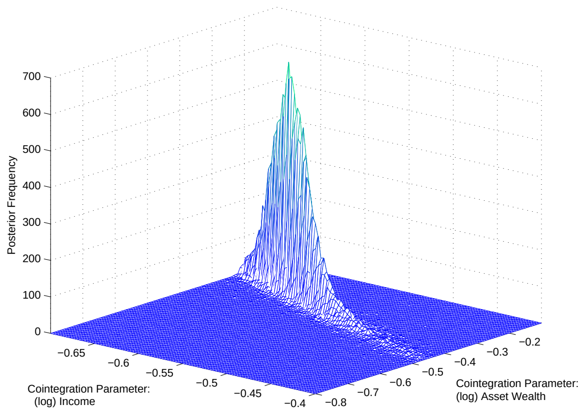

Figure 7: Posterior Draws of Adjustment Parameters

The marginal posterior $p(\alpha_c|c, a, y)$ estimate: sampler, parameter settings and priors are described in the caption to table 3.

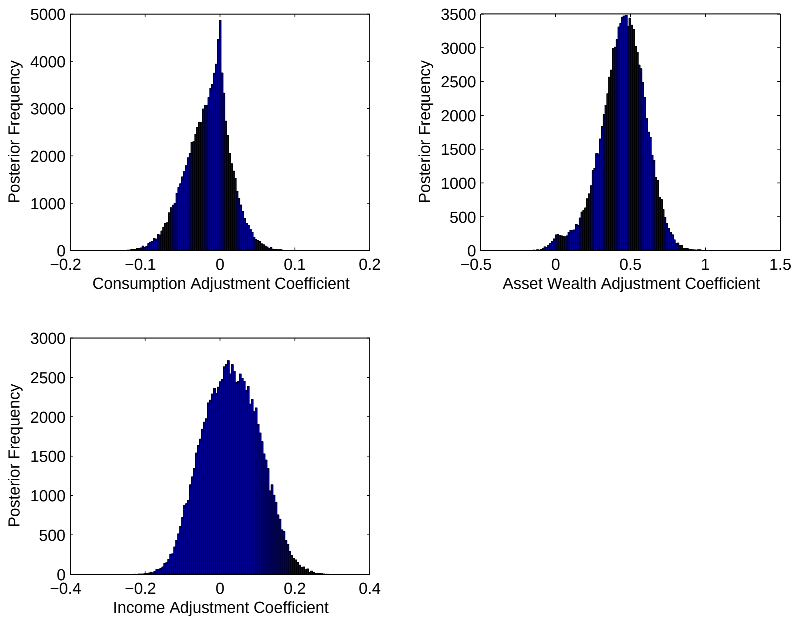

Figure 8: 90% HPD of *cay* Posterior: Q4 1951-Q1 2003

The solid line plots the path of the posterior median value of *cay*. The shaded area is the 90% HPD of the *cay* posterior.

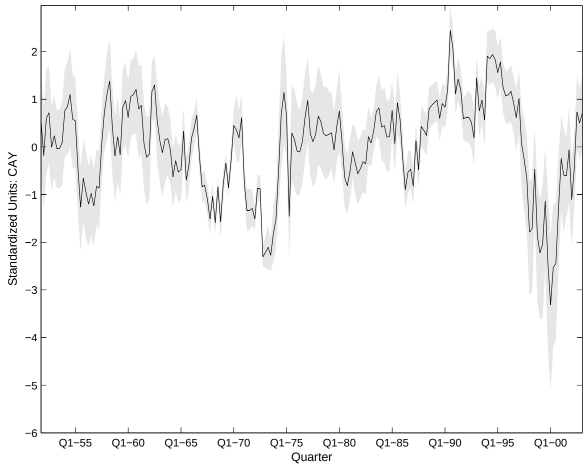

Figure 9: Posterior Distribution of *cay* in Q1 of 2003

The histogram is based on the complete set of the posterior draws obtained from the sampler described in the caption to table 3.

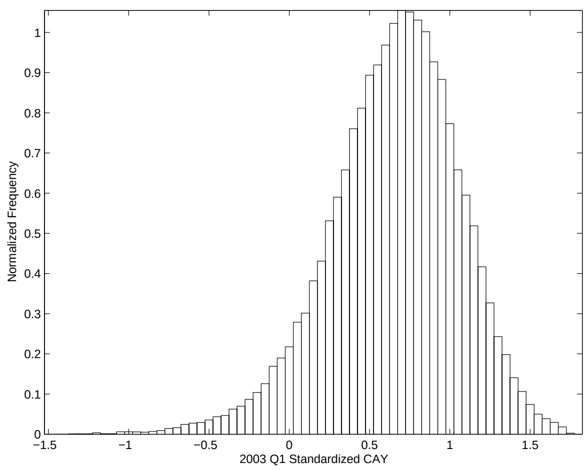

Figure 10: 90% HPD of VAR Implied R-Squared Based on DY, Q4 1951-Q1 2003

The solid line plots the posterior mean of the long horizon R^2 computed using draws from $p(\theta_v|\mathcal{D})$ when DY is used to forecast the quarterly S&P500 return premium. The shaded area denotes the region covered by the 90% HPD interval of the long horizon R^2 implied by the VAR parameters. These results are based on 100,000 sampler draws.

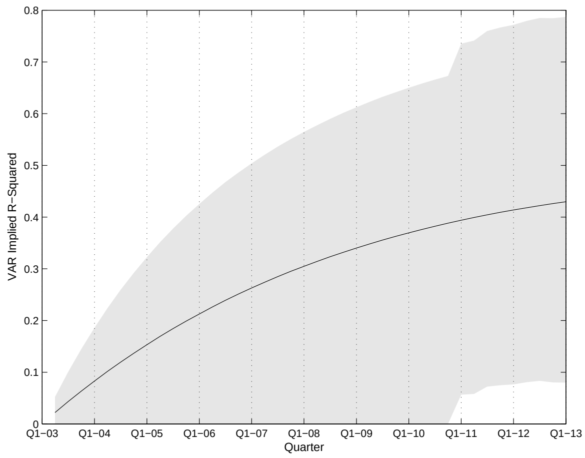

Figure 11: 90% HPD of VAR Implied R-Squared Based on cay

The solid line plots the posterior mean of the long horizon R^2 computed using draws from $p(\theta_v|\theta_0, \mathcal{D})$ when $\hat{cay}$ is used to forecast the quarterly S&P500 return premium. The shaded area denotes the region covered by the 90% HPD interval of the long horizon R^2 implied by the VAR parameters. The dashed lines map the limits of the 90% HPD interval based on $p(\theta_v|\mathcal{D})$ when cay is used to predict the premium. The sampling scheme used to compute $p(\theta_v|\mathcal{D})$ in this case is detailed in Appendix I.

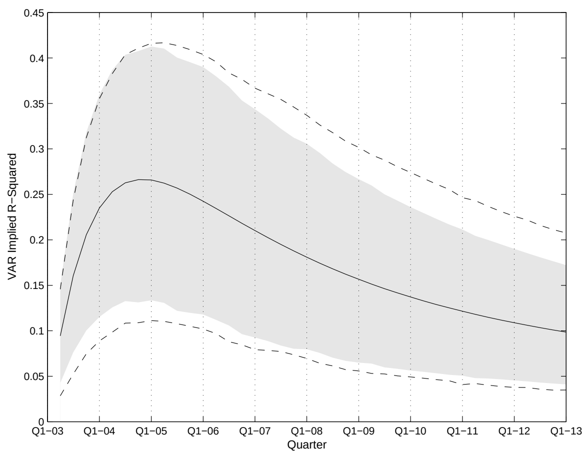

Figure 12: Equity Allocation, Conditional on c, a and y.

The solid line plots equity allocation by horizon conditional on c_{t-1}, a_{t-1}, and y_{t-1} being set equal to their values observed by an investor in Q3, 1953. The investor's risk aversion parameter is 5. The solid line with pentagrams plots the corresponding allocations of a power utility investor with a risk aversion parameter of 20. The dashed line plots equity allocation by horizon conditional on c_{t-1}, a_{t-1}, and y_{t-1} being set equal to their values observed by an investor in Q3, 2002. The investor's risk aversion parameter is 5. The dashed line with pentagrams plots the corresponding allocations of a power utility investor with a risk aversion parameter of 20.

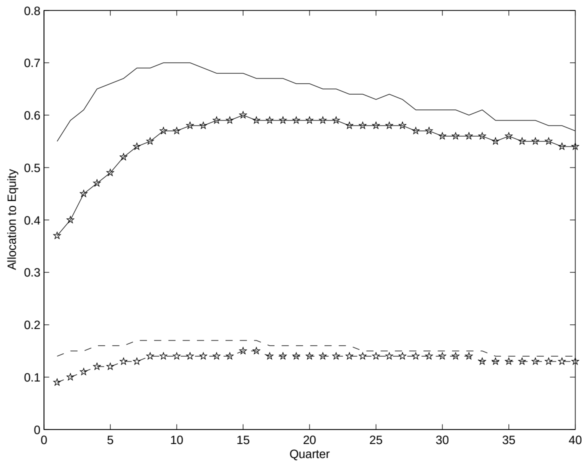

Figure 13: Consumer Sentiment, Predictability and Equity Allocation

The solid line plots equity allocation when the current level of $\hat{cay}$ is one standard deviation above its historical mean, and the ICS is one standard deviation above its historical mean. The dashed line plots equity allocation when the current value of $\hat{cay}$ is one standard deviation above its historical mean, and the ICS is one standard deviation below its historical mean. The dotted line plots equity allocation when the current value of both $\hat{cay}$ and the ICS are set to one standard deviation below historical means. Dash-dots plot equity allocation when the current value of $\hat{cay}$ is a standard deviation down and ICS is a standard deviation up. In each case, the investor's preferences are described by a power utility function with a risk aversion coefficient of 10. All results are based on 100,000 sampler draws. The results in this figure condition on the posterior mean of the cointegrating parameters but account for estimation risk in the (tri-variate) VAR parameters.

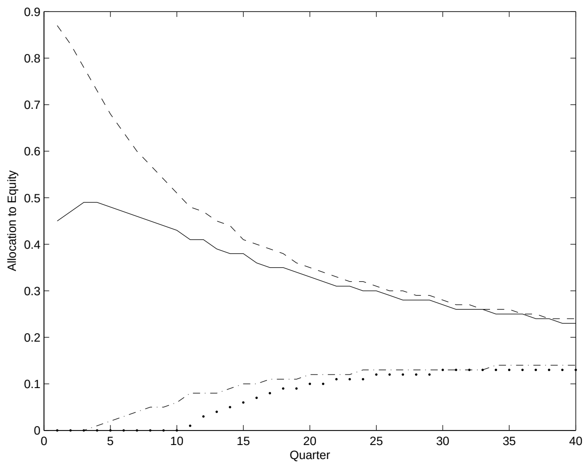

Table 1: Predictive Regressions: S&P 500 Excess Return, Q1 1960-Q1 2003

This table reports predictive regressions of the form $r_t = \hat{\lambda}_0 + \hat{\lambda}_1 pred_{1,t-1} + \ldots + \hat{\lambda}_k pred_{k,t-1}$, where $pred_{j,t-1}$ are lagged predictors. In the $\hat{cay}$ regression: $\hat{cay} = c - 0.2769a - 0.6116y$. Point estimates are based on the posterior mean of the cointegrating vector. For details of $\hat{cay}$ estimation refer to section I. A single star indicates significance of t-statistic at the 5% level, two stars indicate significance at the 2% level, three stars indicate significance at 1%.

Predictor(s)	$\hat{\lambda}_0$	$\hat{\lambda}_1$	$\hat{\lambda}_2$	$\hat{\lambda}_3$	R^2, $(\bar{R}^2)$
$\hat{cay}$	0.0132	0.0261			0.085, (0.0787)
	(1.85)	(3.67***)			
DY	-0.019	0.012			0.0178, (0.011)
	(-0.89)	(1.62)			
ICS	0.0132	-0.008			0.0086, (0.0017)
	(1.78)	(-1.12)			
$\hat{cay}$ [1], DY [2]	0.005	0.025	0.0031		0.086, (0.0734)
	(0.231)	(3.28***)	(0.4)		
$\hat{cay}$ [1], ICS [2]	0.0222	0.0258	0.0013		0.085, (0.0725)
	(0.417)	(3.47)	(-0.171)		
$\hat{cay}$ [1], ICS [2], $ICS \times \hat{cay}$ [3]	0.0087	0.032	-0.0043	-0.0167	0.0167, (0.104)
	(1.197)	(4.144***)	(-0.58)	(-2.46**)	
DY [1], ICS [2], $ICS \times DY$ [3]	-0.025	0.0125	0.0196	-0.0066	0.025, (0.0049)
	(-0.8599)	(1.223)	(0.919)	(-1.05)	
$\hat{cay}$ [1], $ICS \times \hat{cay}$ [2]	0.0088	0.0329	-0.016		0.12, (0.108)
	(1.22)	(4.36***)	(-2.4**)		
DY [1], $ICS \times DY$ [2]	-0.0118	0.0089	-0.0015		0.0196, (0.006)
	(-0.4687)	(0.9443)	(-0.5105)		

Table 2: Predictive VAR Parameter Estimates, Q4 1951-Q1 2003

This table reports the posterior mean of VAR parameters, as per equation (6). The σ_{ij} elements of the residual covariance matrix Σ are reported as correlations for all $i \neq j$. All results pertaining to $p(b_{ij}, \Sigma|data)$ are based on 100,000 draws from the Gibbs sampler.

Equation	Predictor	b_{i0}	b_{i1}	σ_{i1}	σ_{i2}
$[r_t - r_{ft}]$	$c\hat{a}y_{t-1}$	0.0175	2.0251	0.0063	**-0.5784**
$c\hat{a}y_t$	$c\hat{a}y_{t-1}$	0	0.835		0.4805e-005
$[r_t - r_{ft}]$	$[DY_{t-1}]$	0.0854	0.03	0.0068	**-0.6925**
$[DY_t]$	$[DY_{t-1}]$	-0.0325	0.9883		0.004
Monthly Estimates fromBarberis (2000)					
$[r_t - r_{ft}]$	$[DY_{t-1}]$	-0.0143	0.5118	0.0017	**-0.9351**
$[DY_t]$	$[DY_{t-1}]$	0.0008	0.9774		3.0e-006

Table 3: Descriptive Statistics based on Joint Posterior of Cointegrating Parameters and Adjustment Coeffcients, Q4 1951-Q1 2003

The posterior $p(\alpha, \beta, \Gamma, \Phi, \Omega | c, a, y)$ estimates are based on 100,000 draws of a Metropolis-Hastings sampling scheme with a rejection rate of 20.8%. The cointegrating vector $\beta = [\ \beta_c \quad \beta_a \quad \beta_y\ ]'$, and adjustment coefficient vector $\alpha = [\ \alpha_c \quad \alpha_a \quad \alpha_y\ ]'$ component subscripts denote coefficients for (log): consumption, asset wealth and income. The cointegrating vector β incorporates the normalization $\beta_c = 1$. In terms of equation (I.3) $k = 4$ and $w = 1$, that is, each equation in the system (I.3) includes 5 lagged differences of the dependent variable and an intercept. The posterior results are based on diffuse priors as detailed in section I. The reported results are conditional on $k = 6$ and $r = 1$.

	Cointegrating Vector		Adjustment Coefficients		
Parameter	β_a	β_y	α_c	α_a	α_y
mean	-0.2766	-0.6119	-0.0344	0.4671	-0.0529
median	-0.2727	-0.616	-0.0325	0.476	-0.0563
St. Dev.	0.0481	0.0526	0.0342	0.167	0.0728