DISCUSSION PAPER SERIES

No. 3105

MEASURING AND MODELLING VARIATION IN THE RISK-RETURN TRADE-OFF

Martin Lettau and Sydney Ludvigson

FINANCIAL ECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP3105.asp**

ISSN 0265-8003

MEASURING AND MODELLING VARIATION IN THE RISK-RETURN TRADE-OFF

Martin Lettau, Federal Reserve Bank of New York and Stern School of Business, New York University and CEPR
Sydney Ludvigson, New York University

Discussion Paper No. 3105
December 2001

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 3105

December 2001

ABSTRACT

Measuring and Modelling Variation in the Risk-Return Trade-off*

Are excess stock market returns predictable over time and, if so, at what horizons and with which economic indicators? Can stock return predictability be explained by changes in stock market volatility? How does the mean return per unit risk change over time? This chapter reviews what is known about the time-series evolution of the risk-return trade-off for stock market investment, and presents some new empirical evidence using a proxy for the log consumption-aggregate wealth ratio as a predictor of both the mean and volatility of excess stock market returns. We characterize the risk-return trade-off as the conditional expected excess return on a broad stock market index divided by its conditional standard deviation, a quantity commonly known as the Sharpe ratio. Our own investigation suggests that variation in the equity risk-premium is strongly negatively linked to variation in market volatility, at odds with leading asset pricing models. Since the conditional volatility and conditional mean move in opposite directions, the degree of countercyclicality in the Sharpe ratio that we document here is far more dramatic than that produced by existing equilibrium models of financial market behaviour, which completely miss the sheer magnitude of variation in the price of stock market risk.

JEL Classification: G10 and G12
Keywords: consumption, expected returns, Sharpe ratio and volatility

Martin Lettau
Stern School of Business
New York University
44 West Fourth Street, Suite 9-190
New York, NY 10012-1126
USA
Tel: (1 212) 998 0378
Fax: (1 212) 995 4233
Email: mlettau@stern.nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=128074

Sydney Ludvigson
Department of Economics
New York University
269 Mercer Street, 7th floor
New York, NY 10003
USA
Tel: (1 212) 998 8927
Fax: (1 212) 995 4186
Email: sydney.ludvigson@nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=147057

* This Paper has been prepared for the *Handbook of Financial Econometrics*, edited by Yacine Ait-Sahalia and Lars Peter Hansen. Updated versions, along with the monthly and quarterly data, may be found at http://www.stern.nyu.edu/~mlettau and http://www.econ.nyu.edu/user/ludvigsons/. Nathan Barczi and Adam Kolasinski provided excellent research assistance. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Submitted 20 November 2001

1 Introduction

Financial markets are often hard to understand. Stock prices frequently seem volatile and unpredictable, and researchers have devoted significant resources to understanding the behavior of expected returns relative to the risk of stock market investment. Are excess stock market returns predictable over time and, if so, at what horizons and with which economic indicators? Can stock return predictability be explained by changes in stock market volatility? How does the mean return per unit risk change over time? For academic researchers, the progression of empirical evidence aimed at these questions has presented a continuing challenge to asset pricing theory and an important roadmap for future inquiry. For many investment professionals, finding practical answers to these questions is the fundamental purpose of financial economics, as well as its principal reward.

Despite both the theoretical and practical importance of these issues, relatively little is known about how the risk-return tradeoff varies over the business cycle or with key macroeconomic indicators. This chapter reviews the state of knowledge on such variation for stock market investment, and presents some new empirical evidence based on information contained in aggregate consumption and aggregate labor income. We define the risk-return tradeoff as the conditional expected excess return on a broad stock market index divided by its conditional standard deviation, a quantity commonly known as the *Sharpe ratio*. Our study focuses not on the unconditional value of this ratio, but on its evolution through time.

Understanding the time-series properties of the Sharpe ratio is crucial to the development of theoretical models capable of explaining observed patterns of stock market predictability and volatility. For example, Hansen and Jagannathan (1991) showed that the maximum value of the Sharpe ratio places restrictions on the volatility of the set of discount factors that can be used to price returns. The same basic reasoning implies that the pattern of time-series variation in the Sharpe ratio will also place restrictions on the set of discount factors capable of pricing equity returns. The behavior of the Sharpe ratio over time is also fundamental for assessing whether stocks are safer in the long run than they are in the short run, as is increasingly advocated by popular guides to investment strategy (e.g., Siegel (1998)). Only if the Sharpe ratio grows more quickly than the square root of the horizon–so that the variance of the return grows more slowly than its mean–are stocks safer investments in the long run than they are in the short run. Such a dynamic pattern is not possible if stock returns are unpredictable, i.i.d. random variables. Thus, understanding the time-series behavior of the Sharpe ratio not only provides a benchmark for theoretical progress, it has profound implications for investment professionals concerned with strategic asset allocation.

The two components of the risk-return relation (the numerator and the denominator of the Sharpe ratio) are the conditional mean excess stock return, and the conditional standard deviation of the excess return. We focus here on empirically measuring and statistically modeling each of these components separately, a process that can be unified to reveal an estimate of the conditional Sharpe ratio, or *price* of stock market risk. Section 2 discusses estimation of the conditional mean of excess stock returns. In this section we evaluate the statistical evidence for stock return predictability and review the range of indicators with which such predictability has been associated. Taken together, this evidence suggests that excess returns on broad stock market indexes are predictable at long-horizons, implying that the reward for bearing risk varies over time.

One possible explanation for time-variation in the equity risk premium is time variation in stock market volatility. Section 3 reviews the evidence for time-variation in stock market volatility. In many classic asset pricing models, the equity risk premium varies proportionally with stock market volatility. These models require that periods of high excess stock returns coincide with periods of high stock market volatility, implying a constant price of risk. It follows that variation in the equity risk premium must be perfectly positively correlated with variation in stock market volatility.

The important empirical question is whether such a positive correlation between the mean and volatility of returns exists, implying a constant Sharpe ratio. Section 4 ties the evidence on the conditional mean of excess returns in with that on the conditional variance to derive implications for the time-series behavior of the conditional Sharpe ratio. Existing empirical evidence on the sign of the relationship between the conditional mean and the conditional volatility of excess stock returns is mixed and somewhat weak. This may be because some studies have relied on parametric or semi-parametric ARCH-like models of volatility that impose a relatively high degree of structure about which there is little direct empirical evidence. Others have used predictive variables for volatility that are only weakly related to the first moments of returns, and vice versa. Finally, it has been difficult to explain high risk premia with high volatility because evidence suggests that returns are predicable at quarterly and longer horizons, while variation in stock market volatility has, to date, been most evident in high frequency (e.g., daily) data.

In addition to reviewing existing evidence, this chapter presents some new evidence on the risk-return tradeoff. We find that a proxy for the log consumption-aggregate wealth ratio, a variable shown elsewhere to predict excess returns and constructed using information on aggregate consumption and labor income, is also a strong predictor of stock market volatility. To the best of our knowledge, this study is the first to find evidence that at least one variable that strongly forecasts mean returns also strongly forecasts the volatility of returns. Moreover, these results show that the evidence for changing stock market risk is not confined to high frequency data: stock market volatility is forecastable over horizons ranging from one month to six years.

These findings imply that movements in the equity risk-premium are indeed linked empirically to stock market volatility. In addition, the predictability patterns we find for excess returns and volatility imply pronounced countercyclical variation in the Sharpe ratio. This evidence weighs against many time-honored asset pricing models that specify a constant price of risk (for example, the static capital asset pricing model (CAPM) of Sharpe (1964) and Lintner (1965)), and toward more recent paradigms capable of rationalizing a countercyclical Sharpe ratio (e.g., Campbell and Cochrane (1999); Barberis, Huang, and Santos (2001)). Although these more recent frameworks imply a positive correlation between the conditional mean and conditional volatility of returns, unlike the classic asset pricing models this positive correlation is not perfect, making countercyclical variation in the Sharpe ratio possible.

Yet despite evidence that the Sharpe ratio varies countercyclically, our results nevertheless present a problem for modern-day asset pricing theory. Instead of finding a *positive* correlation between the condition mean and conditional volatility, we find a strong *negative* correlation, consistent with the findings of a number of other studies discussed below. More significantly, since the conditional volatility and conditional mean move in opposite directions, the *magnitude* of countercyclicality in the Sharpe ratio that we document here is far more dramatic than that produced by leading asset pricing models capable of generating a countercyclical price of risk. These results suggest that predictability of excess stock returns cannot be readily explained by changes in stock market volatility.

Even if stock market volatility were constant, predictable variation in excess stock returns might be explained by time variation in *consumption* volatility. In a wide-range of equilibrium asset pricing models, more risky consumption streams require asset markets that, in equilibrium, deliver a higher mean return per unit risk. Some variation in aggregate consumption volatility is evident in the data, as we document here. However, this variation is small and we conclude that changes in consumption risk, as measured by changes in the volatility of consumption growth, are insufficiently important empirically to explain the extreme swings in the Sharpe ratio that we find here. These findings imply that even our best-fitting asset pricing models completely miss the sheer magnitude of volatility in the risk-return tradeoff, leaving a "Sharpe ratio volatility puzzle"that remains to be explained. We discuss these issues further in Section 4. Section 5 provides a summary and concluding remarks.

Throughout this chapter, as we consider the evidence for predictability in asset markets, we stress a recurrent theme: the importance of real macroeconomic indicators for estimating the risk-return tradeoff of broad stock market indexes. In particular, we emphasize the usefulness of cointegration–between measures of asset market value and key macroeconomic variables–for understanding these patterns. Such a reliance on macroeconomic data is not common in empirical finance (and might be regarded by some as unorthodox, if not heretical). The relative obscurity of this approach may be partly attributable to the fact that macroeconomic data are subject to a number of measurement limitations not shared by financial market data, and partly because, until recently, empirical connections between the real and financial sectors of the economy have proven difficult to uncover. As a result, research in financial economics has often proceeded independently of that in macroeconomics. Indeed, some researchers have famously mused that the stock market may be little more than a sideshow for macroeconomic activity.[1]

This chapter nevertheless underscores empirical evidence which suggests that the future path of both equity returns and stock market volatility can be usefully informed by observations on real macroeconomic variables. Although the stock market may be a side show for real activity in the short run, the logic of a simple household budget constraint implies that macroeconomic aggregates such as consumption and labor income are inextricably tied to asset values in the long run. Once this long-run equilibrium relation has been identified, deviations from it can be exploited to address questions about the short-run dynamics of asset returns. We use such an approach here to estimate how the risk-return tradeoff on broad stock market indexes evolves over time.

[1] For example, Shleifer (1995).

2 The Conditional Mean of Stock Returns

We capture the risk-return tradeoff for a broad stock market return, R_{st}, by its conditional Sharpe ratio, defined

$$SR_t \equiv \frac{E_t(R_{st+1}) - R_{ft}}{E_t V_{t+1}}, \tag{1}$$

where $E_t(R_{st+1})$ is the mean net return from period t to period $t+1$, conditional on information available at time t; R_{ft}, the risk-free rate, is a short term interest rate paying a return from t to $t+1$, which is known at time t . Similarly, $E_t V_{t+1}$ is a measure of the volatility of the excess return, defined as the standard deviation, conditional on information available at time t. The Sharpe ratio is an intuitively appealing characterization of aggregate stock market returns. It measures how much return an investor can get per unit of volatility in the asset.

The numerator of the Sharpe ratio is the conditional mean excess return. If excess stock returns are predictable, this mean moves over time. The early empirical literature on predictability generally concluded that stock returns were unforecastable, but research in the last 20 years has found compelling evidence of predictability in stock returns. In addition, an active area of recent theoretical research has shown that such predictability is not necessarily inconsistent with market efficiency: forecastablity of equity returns can be generated by time-variation in the rate at which rational, utility maximizing investors discount expected future income from risky assets. Prominent theoretical examples in this tradition include models with time-varying risk aversion (e.g., Campbell and Cochrane (1999)), and models with idiosyncratic risk (e.g., Constantinides and Duffie (1996)).

The evidence for predictability of stock returns has its origins in the literature on stock market volatility. LeRoy and Porter (1981) and Shiller (1981) argued that stock returns were too volatile to be accounted for by variation in future dividend growth, an empirical finding that provides indirect evidence of stock return forecastability. This point may be easily understood by considering an approximate present-value relation for stock market returns. Let d_t and p_t be the log dividend and log price, respectively, of the stock market portfolio, and let $r_{st} \equiv \log(1 + R_{st})$. Throughout this chapter we use lowercase letters to denote log variables, e.g., $\log D_t \equiv d_t$. Campbell and Shiller (1988) show that an approximate expression for the log dividend-price ratio may be written

$$p_t - d_t \approx \kappa + E_t \sum_{j=1}^{\infty} \rho_s^j (\Delta d_{t+j} - r_{s,t+j}), \tag{2}$$

where $\rho_s = P/(P + D)$ and κ is a constant that plays no role in our analysis. This equation is often referred to as the "dynamic dividend growth model" and is derived by taking a first-order Taylor approximation of the equation defining the log stock return, $r_{st} = \log(P_t + D_t) - \log(P_t)$. By taking expectations as of time t (which generates a consistency condition for expectations), the equation says that when the dividend-price ratio is high, agents must be expecting either high returns on assets in the future or low dividend growth rates. If discount rates are constant, expected returns are constant and variation in the price-dividend ratio can only be generated by variation in expected future dividend growth.

How does the volatility of the stock market relate to the forecastability of returns? The answer is immediately evident from (2). Since (2) is derived from an identity, it holds both *ex post* and *ex ante*. Consequently, (2) implies that if discount rates are constant, all of the variability in the price-dividend ratio must be accounted for variability in future dividend growth. Implicitly, LeRoy and Porter (1981) and Shiller (1981) argued that future dividend growth was insufficiently volatile to account for the observed degree of variability in $p_t - d_t$. Equation (2) shows that such a finding implies that forecasts of returns must be time-varying and covary with the dividend-price ratio. Campbell (1991) and Cochrane (1991a) explicitly test this implication and conclude that nearly all the variation in $p_t - d_t$ is attributable not to variation in expected future dividend growth, but to changing forecasts of excess returns.

Equation (2) also demonstrates an important statistical property that is useful for understanding the possibility of predictablility in asset returns. Under the maintained hypothesis that dividend growth and returns follow covariance stationary processes, equation (2) says that the price-dividend ratio on the left-hand-side must also be covariance stationary, implying that dividends and prices are cointegrated. Thus, prices and dividends cannot wonder arbitrarily far from one another, so that deviations of $p_t - d_t$ from its unconditional mean must eventually be eliminated by either, a subsequent movement in dividend growth, a subsequent movement in returns, or some combination of the two. Put another way, cointegration implies that, if the dividend-price ratio varies at all, it must forecast either future returns to equity or future dividend growth, or both. We discuss this property of cointegrated variables further below.

Note that the equity return can always be expressed as the sum of the excess return over a risk-free rate, plus the risk-free rate. It follows that, in principle, variation in the price-dividend ratio could be entirely explained by variability in the expected risk-free rate, even if expected dividend growth rates and risk-premia are constant. In fact, such a scenario is not supported by empirical evidence: variation in expected real interest rates is far too small to account for the volatility of price-dividend ratios on aggregate stock market indexes. Instead, variation in price-dividend ratios is dominated by variation in the reward for bearing risk.[2]

In summary, the early literature on stock market volatility concluded that price-dividend ratios were too volatile to be accounted for by variation in future dividend growth or interest rates alone, thereby providing indirect evidence that expected excess stock returns must vary. A more direct way of testing whether expected returns are time-varying is to explicitly forecast excess returns with some predetermined conditioning variables. For example, equation (2) implies that the price-dividend ratio should provide a rational forecast of long-horizon returns and/or long horizon dividend growth. The empirical asset pricing literature has produced a number of such variables that have been shown, in one subsample of the data or another, to contain predictive power for excess stock returns. Shiller (1981), Fama and French (1988), Campbell and Shiller (1988), Campbell (1991), and Hodrick (1992) find that the ratios of price to dividends or earnings have predictive power for excess returns. Lamont (1998) argues that the dividend payout ratio should be a potentially potent predictor of excess returns, a result of the fact that high dividends typically forecast high returns whereas high earnings typically forecast low returns. Campbell (1991) and Hodrick (1992) find that the relative T-bill rate (the 30-day T-bill rate minus its 12-month moving average) predicts returns, and Fama and French (1989) study the forecasting power of the term spread (the 10-year Treasury bond yield minus the one-year Treasury bond yield) and the default

[2] See Campbell, Lo, and MacKinlay (1997), chapter 8 for summary evidence.

spread (the difference between the BAA and AAA corporate bond rates). We denote these last three variables $RREL_t$, TRM_t, and DEF_t respectively. Finally, Lewellen (1999) and Vuolteenaho (2000) forecast returns with an aggregate book-market ratio. Various methodologies for forecasting returns have been employed, including in-sample forecasts based on direct regressions of long-horizon returns on predictive variables, vector autoregressive approaches which impute long-horizon statistics rather than estimating them directly, and a battery of out-of-sample procedures aimed at testing for subsample stability and overcoming small sample biases in statistical inference. We discuss these procedures further below.

It is commonly believed that expected excess returns on common stocks vary with the business cycle, so that risk-premia are higher in recessions than they are in expansions. If so, excess stock returns should be forecastable by business cycle variables at cyclical frequencies. The difficulty with this belief is that the financial variables used as forecasting indicators in the studies cited above have been most successful at predicting returns over long horizons. Over horizons spanning the length of a typical business cycle, stock returns have typically been found to be only weakly forecastable by these variables. Moreover, until recently, cyclical macroeconomic variables appeared to display very little if any relation to predictable variation in stock returns. We now discuss an empirical investigation that takes a new approach to investigating the linkages between the real macroeconomy and financial markets.

Lettau and Ludvigson (2001a) study the forecasting power for stock returns of a proxy for the log consumption-aggregate wealth ratio, where aggregate wealth, W_t, equals the sum of human capital, H_t and nonhuman capital, or asset wealth, A_t. Lettau and Ludvigson argue that a standard budget constraint identity implies that log consumption, c_t, log labor income, y_t and log nonhuman, or asset, wealth, a_t are cointegrated, and that deviations from the common trend in these variables can be thought of as fluctuations in log consumption-aggregate wealth ratio a variable that is likely to forecast stock returns. Although human capital, H_t is unobservable, they assume that aggregate labor income is nonstationary and defines the trend in human capital, so that the log of human capital may be written, $h_t = y_t + z_t$, where z_t is a stationary random variable. Using this relation, they derive an equation taking the form

$$cay_t \equiv c_t - \alpha_a a_t - \alpha_y y_t \approx k + E_t \sum_{i=1}^{\infty} \rho_w^i \left(r_{w,t+i} - \Delta c_{t+i} \right) + (1 - \omega) z_t, \tag{3}$$

where r_w is the return to aggregate wealth, ρ_w is the steady-state ratio of new investment to total wealth, $(W - C)/W$, and k is a constant that plays no role in our analysis. Equation (3) is derived by taking a Taylor expansion of the equation defining the evolution of aggregate wealth, $W_{t+1} = (1 + R_{w,t+1})(W_t - C_t)$.

Several points about equation (3) bear noting. First, under the maintained hypothesis that returns, consumption growth and labor income growth are stationary, the left-hand-side of (3) (denoted cay_t for short) is observable as a cointegrating residual for consumption, asset wealth and labor income. The equation says that, while asset values may deviate from consumption and labor income in the short-run, they are tied to these variables in the long-run. It follows that knowing where asset values are in relation to consumption and labor income may provide valuable information about the future path of wealth. As we will see, this intuition lies at the heart of the forecastability of U.S. excess stock market returns.

Second, the parameters of this cointegrating relation in principle give steady state wealth shares, with α_a equal to the average share of asset wealth in aggregate wealth and α_y equal to the average share of human capital in aggregate wealth. In practice, data measurement considerations for aggregate consumption imply that these coefficients are likely to sum to a number less than one since only a subset of aggregate consumption based on nondurables and services is used (see Lettau and Ludvigson (2001a)). Third, although cay_t is proportional to $c_t - w_t$ only if the last term on the right-hand-side of (3) is constant, Lettau and Ludvigson (2001d) show that this term is primarily a function of expected future labor income growth, which does not appear to vary much in aggregate data. Thus, cay_t may be thought of as a proxy for the log consumption-aggregate wealth ratio, $c_t - w_t$, and we will often refer to it as such in this chapter. Finally, note that stock returns, r_{st}, are but one component of the return to aggregate wealth, $r_{w,t}$. Stock returns, in turn, are the sum of excess stock returns and the real interest rate. Therefore equation (3) says that the log consumption-aggregate wealth ratio embodies rational forecasts of excess stock returns, interest rates, returns to nonstock market wealth, and consumption growth, conditional on the expected values of these variables varying over time.

The principal of cointegration is as important for understanding (3) as it is for understanding (2). In direct analogy to (2), (3) says that if the consumption-wealth ratio varies at all, it must forecast either future returns, future consumption growth, future labor income growth (embedded in z_t), or some combination of all three. Thus, cay_t is a possible forecasting variable of stock returns and consumption growth for the same reasons the price-dividend ratio is a possible forecasting variable for stock returns and dividend growth. In contrast to the price-dividend ratio, however, cay_t contains cointegrating parameters that must be estimated, a task that is straightforward using procedures developed by Johansen (1988) or Stock and Watson (1993). Lettau and Ludvigson (2001a) describe these procedures in more detail and apply them to data on aggregate consumption, labor income and asset wealth to obtain an estimate of cay_t, denoted $\widehat{cay}_t$. Using an updated sample spanning the period from the fourth quarter of 1952 to the first quarter of 2001, we estimate a value here for $cay_t = c_t - 0.61 - 0.30a_t - 0.60y_t$.[3]

We now have a number of variables, based on both financial and macroeconomic indicators, that have been documented, in one study or another, to predict excess stock market returns. How does the predictive capacity of these forecasting variables compare? To summarize the empirical findings of this literature, Table 1 presents the results of in-sample predictive regressions of quarterly excess returns on the value-weighted index provided by the Center for Research in Securities Prices (CRSP-VW), in excess of the return on a three-

[3]As in Lettau and Ludvigson (2001a), we use nondurables and services expenditure as our measure of c_t in cay_t. Since nondurables and services expenditures are typically about 85 percent of total personal consumption expenditures whereas the budget constraint identity applies to total expenditures, we showed there that the cointegrating coefficients obtained using this measure of c_t are likely to sum to a number less than one. Notice however, that if total consumption expenditures were used in place of nondurables and services expenditures, the theory implies that value of these cointegrating coefficients should sum to one. Applying the same procedure when the log of total personal consumption expenditures, c_t^T is used in place of c_t, we obtain the estimated value $\widehat{cay}_t^T = c_t^T - 0.70 - 0.43a_t - 0.58y_t$. The cointegrating coefficients from this estimation are all strongly statistically significant (with t-statistics in excess of 14) and, more important, sum to almost exactly to one. For the empirical investigation of asset returns that we undertake here, however, we continue to use nondurables and services in our construction of cay_t. We do this because both theory and evidence suggest that such a measure of consumption should be closer to a random walk than durables expenditures which are included in c_t^T (Mankiw (1982)). This random walk property of nondurables and services expenditure maximizes the chance that cay_t will have predictive power for future asset returns if expected returns are in fact time-varying, because it facilitates identification of the transitory component in wealth (see Lettau and Ludvigson (2001d)). To understand why, note that, if consumption and labor income are each close to random walk processes and cointegrated with asset values, deviations from the common trend in consumption, asset wealth and labor income (i.e., movements in cay_t) can only be related to future movements in wealth, rather than to some combination of future movements in all three variables. See Lettau and Ludvigson (2001d) for further discussion.

month Treasury bill rate. This table is an updated version of the results presented in Lettau and Ludvigson (2001a), which used data from the fourth quarter of 1952 to the third quarter of 1998. Here we compare the forecasting power of $\widehat{cay}_t$, the dividend-price ratio, $RREL_t$, TRM_t, and DEF_t. Let r_{st} denote the log real return of the CRSP value-weighed index and $r_{f,t}$ the log real return on the 3-month Treasury bill (the 'risk-free' rate). The log excess return is $r_{st} - r_{f,t}$. Log price, p_t, is the natural logarithm of the CRSP-VW index. Log dividends, d, are the natural logarithm of the sum of the past four quarters of dividends per share. We call $d - p$ the dividend yield.

At a one quarter horizon, the only variables that have marginal predictive power in this sample are the consumption-wealth ratio proxy, $\widehat{cay}_t$, and the relative-bill rate, $RREL_t$. The first row of each panel of Table 1 shows that a regression of returns on one lag of the dependent variable displays no forecastability. By contrast, $\widehat{cay}_t$ explains a substantial fraction of the variation in next quarter's return on the CRSP-VW index. Adding last quarter's value of $\widehat{cay}_t$ to the model allows the regression to predict an additional seven percent of the variation in next period's excess return and an extra six percent of the variation in next period's real return. Panel C of Table 1 shows that neither the dividend yield, the default premium, or the term premium display marginal predictive power for quarterly excess returns. The relative bill rate does, but it adds less than two percent to the adjusted R^2 (compare rows 8 and 9 of Panel C). In short, except for $\widehat{cay}_t$, few variables display important predictive power for returns at quarterly horizons.

Still, the theory behind (2 and 3) makes clear that both the dividend-price ratio and the consumption-wealth ratio should track longer-term tendencies in asset markets rather than provide accurate short-term forecasts of booms or crashes. Thus, Table 2, Panel A, presents the results of long-horizon forecasting regressions of excess returns on the CRSP-VW index, on several lagged forecasting variables. The dependent variable is the H-quarter log excess return on the CRSP-VW index, equal to $r_{t+1} - r_{f,t+1} + ... + r_{t+H} - r_{f,t+H}$. For each regression, the table reports the estimated coefficient on the included explanatory variable(s), the adjusted R^2 statistic, and the Newey-West corrected t-statistic for the hypothesis that the coefficient is zero. (Results using $RREL_t$, TRM_t, and DEF_t as predictive variables indicated that these variables displayed no forecasting power at any horizon in our sample. As a result, those regressions are therefore omitted from the table to conserve space.)

The first row of Table 2 shows that $\widehat{cay}_t$ has significant forecasting power for future excess returns at horizons ranging from one to 24 quarters.[4] The t-statistics are above 3 for all horizons. The predictive power of $\widehat{cay}_t$ is hump-shaped and peaks around three years in this sample; using this single variable alone achieves an $\overline{R}^2$ of 0.26 for excess returns over an 24 quarter horizon. These results provide strong evidence that the conditional mean of excess stock returns varies in U.S. data over horizons of several years.

The remaining rows of Panel A give an indication of the predictive power of other variables for long-horizon excess returns. Row 2 reports long-horizon regressions using the dividend-yield as the sole forecasting variable. These results are quite different to those obtained elsewhere (for example, Fama and French (1988); Lamont (1998); Campbell, Lo, and MacKinlay (1997)) because we use more recent data. The dividend-price ratio has no

[4]The t-statistics reported in the tables are all corrected for generalized serial correlation using the Newey-West procedure (Newey and West (1987)). Although not reported here, the marginal predictive power of $\widehat{cay}_t$ and $RREL_t$ for future returns is also robust to the use of an alternative correction for serial correlation, developed by Hodrick (1992), to address the specific serial correlation structure that arises with the use of overlapping data in long-horizon regressions (see Lettau and Ludvigson (2001c)).

ability to forecast excess stock returns at horizons ranging from one to 24 quarters when data after 1995 are included. The last half of the 1990s saw an extraordinary surge in stock prices relative to dividends, weakening the tight link between the dividend-yield and future returns that has been documented in previous samples. Undoubtedly some of this reduction in predictive power is attributable to recent changes in the way dividends and earnings are paid-out and reported. For example, firms have been distributing an increasing fraction of total cash paid to shareholders in the form of stock repurchases. If the data on dividends do not include such repurchases, changes of this type would distort measured dividends and reduce the forecasting power of the dividend-price ratio. By contrast, data on aggregate consumption is largely free of at least these measurement problems. This may partly explain why Lettau and Ludvigson (2001a) find that the consumption-wealth variable has better predictive power for excess stock returns than all of the financial variables listed above in both in-sample and out-of-sample tests, consistent with the results presented here.

Row 3 of Panel A shows that $RREL_t$ has forecasting power that is concentrated at shorter horizons than $\widehat{cay}_t$, with R^2 statistics that peak at 6 quarters. The coefficient estimates are strongly statistically significant, with t-statistics in excess of 3 at one and two quarter horizons, but the variable again explains a smaller fraction of the variability in future returns than does $\widehat{cay}_t$. Row 4 of Table 2 presents the results of forecasting excess returns using a multivariate regression with $\widehat{cay}_t$, $RREL_t$, and $d_t - p_t$ as predictive variables. The results demonstrate the substantial predictability of excess returns; the adjusted R^2 statistics range from 10 to 32 percent for return horizons from one to 24 quarters.

The data used in the analysis just discussed are quarterly. Lettau and Ludvigson (2001a) also used quarterly data because, although data on c_t and y_t are available on a monthly basis, the data for asset wealth, a_t, contained in $\widehat{cay}_t$, are only reported on a quarterly basis. Nevertheless, it is more common in asset pricing applications to use monthly data in computing long-horizon regression statistics. Therefore, in order to facilitate comparison with existing literature, we construct a monthly estimate of wealth from the quarterly values, and thereby producing a monthly series for $\widehat{cay}_t$. Since $\widehat{cay}_t$ is used merely as a forecasting variable for identifying conditional expectations, there is nothing wrong with constructing a monthly estimate of the series as long as no forward-looking information is used to do so. Thus, we construct this estimate in the following way.

First, we divide total asset wealth into two components, stock and nonstock wealth. The definitions of these components are given in the data appendix at the end of this chapter. Noting that the raw quarterly data for wealth is measured at the end of the period, we fill in an end-of-month value for stock market wealth by multiplying the end-of-quarter value of stock wealth by $(1 + g_1)$, where g_1 is the growth rate of the Standard & Poor 500 stock price index over the first month of the quarter.[5] This provides an estimate of stock wealth as of the end of the first month of the quarter. This value is then multiplied by $(1 + g_2)$, where g_2 is the growth rate of the Standard & Poor 500 stock price index over the second month of the quarter, to get a value for stock wealth as of the end of the second month of the quarter. Since the vast majority of variation in stock wealth at monthly and quarterly horizons is generated by fluctuations in capital gains (see Ludvigson and Steindel (1999)),

[5]We use the Standard & Poor index, which is very highly correlated with the CRSP-VW index, to construct this estimate because it is available continuously. By contrast, the CRSP-VW index is published with a one year lag.

this estimate of monthly stock market wealth should be a reasonable one.

Generating a monthly series for the nonstock component of wealth requires a different procedure since there is no monthly price series for this asset stock. Therefore we fill in end-of-month values for nonstock market wealth by multiplying the end-of-quarter value by $(1 + g_3)$, where g_3 is a moving average of the last eight quarters growth rate of nonstock wealth, measured at a monthly rate. The sum of these monthly stock and nonstock wealth estimates provides an estimate of monthly asset wealth, a_t. Note that no forward-looking information is used to construct these monthly estimates. The end-of-month estimate for month m, a_m, uses only information available at the end of month m. The resulting series is lagged to create a beginning-of-period value, and $\widehat{cay}_t$ is then constructed from monthly data on c_t and y_t in the same manner used to construct this variable from quarterly data.

Panel B of Table 2 shows the results of long-horizon forecasting regressions of excess returns using this monthly series for $\widehat{cay}_t$. The monthly sample runs from January 1959 to December 2000. It is immediately evident that a forecasting pattern similar to that using quarterly data is present. The monthly cointegrating residual, $\widehat{cay}_t$, is strongly statistically significant, with all t-statistics above 4 for return horizons ranging from one to 84 months. The R^2 statistics suggest that the variable explains more than 40 percent of the variability in future returns at horizons in excess of 24 months, higher than that found for the same horizon using quarterly data. The higher R^2 statistics using monthly data may arise for two reasons. First, the monthly series for $\widehat{cay}_t$ uses more recent information on asset values than does the quarterly series. Both wealth series are measured at the beginning of the period. Using quarterly data, this implies that forecasting returns in the first quarter of a given year uses information as of the end of September of the previous year. By contrast, the monthly data set uses information as of the end of November of the previous year.

Second, the monthly and quarterly regressions differ in the number of overlapping residuals relative to the number of observations. It is well known that asymptotic theory may be misleading in finite samples when the horizon K is large relative to the sample size (Hodrick (1992); Nelson and Kim (1993); Richardson and Stock (1989)). More importantly, Valkanov (2001) shows that the finite-sample distributions of R^2 statistics in long-horizon regressions do not converge to their population values when there are overlapping residuals, with the degree of divergence dependent on the amount of overlap. In addition, Valkanov shows that the R^2 statistic in the types of long-horizon regressions conducted above behaves like a random variable, implying that different data sets (e.g., quarterly versus monthly) are likely to generate different distributions for this statistic. At the same time, however, Valkanov's results imply that, for either data set, a 90 percent confidence interval would not include values as high as 40 percent for the R^2 statistic under the null of no return predictability. It follows that the results in Table 2 using monthly data on $\widehat{cay}_t$ constitute evidence of long-horizon predictability.

Finite sample problems with overlapping data in long-horizon regressions may be avoided by using vector autoregregressions (VARs) to impute the long-horizon R^2 statistics, rather than estimating them directly from long-horizon returns. This approach assumes that the dynamics of the data may be well described by a VAR of a particular lag order, implying that conditional forecasts over long-horizons follow directly from the VAR model. The methodology for measuring long-horizon statistics by estimating a VAR has been covered by Campbell (1991), Hodrick (1992), and Kandel and Stambaugh (1989), and we refer the reader to those articles for further details. We present the results of using this methodology in Table 3, which investigates the long horizon predictive power of $\widehat{cay}_t$ using a bivariate, first-order VAR for returns and $\widehat{cay}_t$. The top panel presents results using the quarterly data set; the bottom panel using the monthly data set. For each return horizon we consider, we calculate an implied R^2 statistic using the coefficient estimates of the VAR and the estimated covariance matrix of the VAR residuals. Notice that the pattern of the implied R^2 statistics is very similar to those from the produced from the single equation regressions in Table 2. This suggests that evidence favoring predictability in Table 2 is not merely a spurious result arising problems with overlapping data in finite samples. For both monthly and quarterly data, the implied R^2 statistics for forecasting excess stock returns with $\widehat{cay}_t$ are hump-shaped in the horizon and peak around two years. This evidence confirms the findings based on direct long horizon regressions, implying that excess returns contain a predictable component that is concentrated at horizons in excess of one year.

2.1 Statistical Issues With Forecasting Returns

The results presented above indicate that excess equity returns are forecastable, suggesting that equity risk-premia vary with time. There are however a number of statistical issues that arise in applying these forecasting tests. One, discussed above, concerns the biases in finite samples that can arise with the use of overlapping data in finite samples. A second concerns statistical inference when returns are regressed on a persistent, predetermined regressor. For example, Stambaugh (1999) showed that the ordinary least squares (OLS) estimate of the regression coefficient of the type reported in Table 2 is biased up in finite samples when the return innovation covaries with the innovation in the forecasting variable. Furthermore, Stambaugh shows that this bias is increasing in the degree of persistence of the forecasting variable. Ferson, Sarkissian, and Simin (1999) report similar findings. Other researchers have conducted explicit finite samples tests and concluded that evidence of predictability may be weaker than previously thought. Ang and Bekaert (2001) carefully compute empirical t-statistics from Monte Carlo simulations of finite sample data under the null of no predictability and conclude that the dividend-yield has little forecasting power for future returns in monthly data.[6] Still others have found that the dividend yield has no ability to predict out-of-sample despite its ability to do so in-sample (for example, Bossaerts and Hillion (1999); Goyal and Welch (1999)). In summary, these studies show that the researcher may find spurious evidence of return predictability in small samples, or evidence of instability in the predictive relation, especially when the forecasting variable is sufficiently persistent.

While formal statistical evidence of predictability may be weak using indicators such as the dividend-yield or earnings-price ratios as predictive variables, it is strong using $\widehat{cay}_t$ as a predictive variable. Lettau and Ludvigson (2001a) test return forecastability by $\widehat{cay}_t$ using a bootstrap procedure that addresses all of the concerns raised by each of the studies cited above. One possible reason for the superior performance of $\widehat{cay}_t$ is that this indicator has the important advantage of being less persistent than the dividend yield.[7]

[6] It is not clear, however, whether this result for post-war U.S. data on the Standard & Poor 500 index is being driven by the inclusion of data from the last five years. Many other researchers have documented that excess returns on the Standard & Poor 500 index are forecastable using the dividend-yield, in post-war data through 1994 (e.g., Fama and French (1988); Campbell (1991); Hodrick (1992); Campbell, Lo, and MacKinlay (1997); Lamont (1998); Lettau and Ludvigson (2001a).

[7] Lettau and Ludvigson (2001a) report a first order correlation coefficient for $d_t - p_t$ equal to 0.93 in quarterly data, compared to 0.79 for $\widehat{cay}_t$.

Lettau and Ludvigson address the statistical issues raised in the papers cited above by conducting quarterly out-of-sample forecasts of excess returns using $\widehat{cay}_t$. They and find strong evidence of out-of-sample predictability for excess returns on both the CRSP-VW and Standard & Poor 500 stock market indices. Lettau and Ludvigson also ask whether the out-of-sample predictability is statistically significant by employing the two statistical tests found by Clark and McCracken (1999) to have the best overall power and size properties. The first is a modified Harvey, Leybourne, and Newbold (1998) test statistic adapted to address the fact that the limiting distribution of this test statistic is non-normal when the forecasts are nested under the null. This statistic provides a test of the null hypothesis that a constant expected returns model "encompasses" all the relevant information for next period's value of the dependent variable, against the alternative that the $\widehat{cay_t}$ predictive model contains additional information. The second is an out-of-sample F-type test, developed in McCracken (1999), of equal mean-squared forecasting error. The null hypothesis for this test is that a constant expected returns model, which excludes $\widehat{cay_t}$, has a mean-squared forecasting error that is less than or equal to that of a model which includes $\widehat{cay_t}$; the alternative is that the $\widehat{cay_t}$-augmented model has lower mean-squared error. To assess the role of small-sample biases, Lettau and Ludvigson report bootstrapped critical values of these statistics, generated by performing repeated simulations under the null hypothesis of no predictability.[8] This procedure addresses the concerns raised in all of the above studies for other forecasting variables: the bootstrapped statistics are for out-of-sample forecasts and are determined by the data sample used, under the null hypothesis of no return predictability. Thus, the

[8]During each iteration of each simulation, the parameters in $\widehat{cay}_t$ are reestimated.

procedure provides a very stringent statistical hurdle to establishing predictability of excess stock returns by $\widehat{cay}_t$. Lettau and Ludvigson nevertheless find that $\widehat{cay}_t$ has predictive power that passes even these stringent tests: at a quarterly horizon, a model of constant expected returns is rejected in favor of a model of time-varying expected returns when $\widehat{cay}_t$ is used as the predictive variable. These findings provide compelling evidence that excess returns are forecastable and suggest that authors such as Ang and Bekaert (2001) overstate the case for no return predictability based on statistical criteria using more persistent dividend-price ratios or earnings-price ratios as predictive variables.

2.2 Conceptual Issues With Forecasting Returns

Several conceptual issues arise when considering the evidence for time-variation in expected excess stock returns. Consider, for example, using the log dividend-price ratio as a predictor of excess returns. Studies that conduct such an analysis typically assume, either explicitly or implicitly, that the ratio of dividends to prices, D_t/P_t, is covariance stationary. This is a reasonable assumption since it is not sensible that prices could wonder arbitrarily far from measures of fundamental value. This assumption implies that the log price-dividend ratio, $p_t - d_t$, is also covariance stationary and that p_t and d_t are cointegrated with cointegrating vector $(1, -1)'$. Under these assumptions, the Granger representation theorem states that variation in $p_t - d_t$ must be related to either variation in future dividend growth, future returns or both, and this relation should be evident in a univariate regressions of either dividend growth or returns on past values of $d_t - p_t$.[9] This means that expected returns

[9] If dividends and prices are cointegrated with cointegrating vector $(1, -1)$, the Granger representation theorem states that $d_t - p_t$ must forecast either Δp_t, or Δd_t, the log difference of dividend growth. Using

cannot be constant if the price-dividend ratio varies, unless expected dividend growth varies. This implication follows on purely statistical grounds from the presumption that D_t/P_t, is covariance stationary. Thus, evidence that expected returns are constant requires not merely that returns be *unforecastable* by $d_t - p_t$, but also that dividend growth be strongly *forecastable* by $d_t - p_t$, and forecastable by the amount necessary to account for the overall variation in $d_t - p_t$. Although the most stringent statistical tests sometimes suggest that $d_t - p_t$ is a weak and/or unstable univariate predictor of returns, the evidence that $d_t - p_t$ predicts dividend growth in post-war U.S. data is even weaker (Campbell (1991); Cochrane (1991b); Cochrane (1994); Cochrane (1997); Ang and Bekaert (2001); Campbell and Shiller (2001)).[10] A similar finding holds for the consumption-wealth ratio proxy: (3) suggests that $\widehat{cay}_t$ could be a long-horizon forecaster of consumption growth, returns, or both. But Lettau and Ludvigson (2001a) find that, although $\widehat{cay}_t$ is a strong predictor of excess stock returns, it has no forecasting power for future consumption growth at any horizon. In addition, it has no forecasting power for future labor income growth, a possibility suggested by the Granger representation theorem for the system $(c_t, a_t, y_t)'$ (Lettau and Ludvigson (2001d)). Instead, Lettau and Ludvigson (2001d) find that household asset values are tied to consumption and labor income in the long-run, and that short-run deviations from this long-run equilibrium

the approximation, $r_{st} \approx \Delta p_t$, it follows that $d_t - p_t$ must forecast either Δd_t or r_{st} up to a first-order approximation.

[10] All of these studies find that dividend growth is unpredictable by the lagged dividend yield in univariate regressions using post-war U.S. data. All but the fifth find that returns are predictable by the lagged dividend yield. It is possible that expected dividend growth and expected returns are *both* time-varying, and that a correlation between the two makes it difficult to statistically distinguish variation in either in a univariate regression. Ang and Bekaert (2001) do find that a multivariate predictive regression for Δd_t provides some evidence of predictability in dividend growth rates. (It is not clear from their results how quantitatively important such predictable variation is, however.) Still, because dividend growth predictability is not evident in a univariate regression of Δd_{t+1} on $d_t - p_t$, their finding provides indirect evidence that expected returns are also time-varying (see (2)).

are exclusively associated with transitory movements in stock market wealth. These results reinforce the conclusion that expected returns are time-varying, even though some statistical tests using some predictive variables fail to provide evidence of predictability. Moreover, the results raise an important point, namely that statistical analyses of return predictability are, by themselves, insufficient for testing whether expected returns are constant.

A separate set of conceptual issues arises in using macroeconomic variables, such as $\widehat{cay}_t$, to forecast returns. Unlike financial data, macroeconomic data are not available in real time and $\widehat{cay}_t$ itself contains parameters that must be estimated. These attributes raise important questions. Should evidence on predictability in returns be based on tests where the parameters in $\widehat{cay}_t$ are reestimated every period, using only data available at the time of the forecast? To answer, we argue that it is essential to distinguish two questions about return forecastability. The first question–the question of concern in this paper–is "Are expected excess returns time-varying?" The second question, of interest to practitioners, is "Can the predictability of returns be statistically detected and exploited in real time?" The questions are distinct, and the empirical approach taken will depend on the question at hand.

To understand this distinction, note that the equilibrium interpretation of (3) given in Lettau and Ludvigson (2001a) says that the true value of $\widehat{cay}_t$ is in agents' information sets at time t, and should provide a rational forecast of either future returns or future consumption growth or both. Such a prediction would be difficult to test without using our best estimate of the underlying data and parameters in cay_t. For example, consistent estimation of the parameters in cay_t requires a large number of observations, and a reestimation of these parameters is likely to induce significant sampling error during the early estimation recursions when only a small amount of data is available. This would make it more difficult for $\widehat{cay}_t$ to display forecasting power of either returns or consumption growth even if it in fact had forecasting power. Similarly, if the first question is of interest, the use of historical data is preferable to real time data since the finally revised versions of these data presumably provide a better estimate of the true values in investors information sets. Thus, when the first question is of interest, the researcher should give the theory underlying the predictive content of $\widehat{cay}_t$ its best chance of success, by using historical data and the full sample to estimate the parameters in $\widehat{cay}_t$. On the other hand, if the second question is of interest, it is reasonable to attempt to emulate the data restrictions a practitioner would face in real time.

In summary, the discussion above outlines the considerable statistical demands in establishing formal evidence of predictability of excess stock returns. While investigations of whether predictive variables are able to meet these challenges are clearly informative, it is also important to recognize that failure to find formal statistical evidence of predictability is not equivalent to finding evidence against predictability. Those results may say more about the power of our statistical tests and the size of our data sets than about the true degree of time-variation in expected returns. Moreover, in some cases, other sources of evidence and economic reasoning may point toward time-variation in expected returns even in cases where statistical inference does not confirm predictability.

3 The Conditional Volatility of Stock Returns

The denominator of the Sharpe ratio defined above, (1), is the conditional standard deviation of excess returns. Although several papers have investigated the empirical determinants of stock market volatility, few have found real macroeconomic conditions to have a quantitatively important impact on conditional volatility. In a classic paper, Schwert (1989) finds that stock market volatility is higher during recessions than at other times, but he also finds that this recession factor–as with measured volatility for range of macroeconomic time series–plays a small role in explaining the behavior of stock market volatility over time. Thus, existing evidence that stock market risk is related to the real economy is at best mixed.

There is even more disagreement among studies that seek to determine the empirical relation between conditional mean and conditional volatility of stock returns. Bollerslev, Engle, and Wooldridge (1988) and Harvey (1989) find a positive relation, while Campbell (1987), Breen, Glosten, and Jagannathan (1989) Glosten, Jagannathan, and Runkle (1993), Whitelaw (1994) and Brandt and Kang (2001) find a negative relation. French, Schwert, and Stambaugh (1987) find a negative relation between returns and the *unpredictable* component of volatility, a finding they interpret as indirect evidence that *ex ante* volatility is positively related to *ex ante* excess returns; but they do not find evidence of a direct connection between these variables.

Empirical studies of the relation between the conditional mean and volatility of stock returns have been based on only a few estimation methodologies. One of the most popular of these methodologies, used by French, Schwert, and Stambaugh (1987), Breen, Glosten, and Jagannathan (1989), and Glosten, Jagannathan, and Runkle (1993), specifies a general empirical specification relating conditional means to conditional volatility taking the form

$$E[r_{st+1} - r_{ft} \mid Z_t] = \alpha + \beta Var(r_{st+1} - r_{ft} \mid Z_t),$$

where Z_t denotes the information set of investors. This information set can contain predetermined predictive variables, or ex ante measures of volatility inherent in a generalized autoregressive conditional heteroskedasticity (GARCH) model as in French, Schwert, and Stambaugh (1987). In a similar approach, Whitelaw (1994) models the conditional mean and volatility separately as functions of predetermined variables and estimates the conditional correlation between the two from the fitted series based on the coefficient estimates of the separate regressions. Brandt and Kang (2001) use a latent vector autoregressive approach to estimating the relation between the conditional mean and conditional volatility, discussed further below.

None of these studies find compelling evidence that variables which forecast means also forecast variances. This may be because some studies have relied on parametric or semi-parametric ARCH-like models that impose a relatively high degree of structure about which there is little direct empirical evidence. Others have used predictive variables for volatility that are only weakly related to the first moments of returns, and vice versa.

The use of $\widehat{cay}_t$ as a predictive variable provides a fresh opportunity to revisit the relation between the conditional first and second moment of returns, for several reasons. First, evidence in Lettau and Ludvigson (2001a) suggests that, of existing predictive variables, $\widehat{cay}_t$ displays the strongest statistical connection to future excess returns on broad stock market indexes. It follows that movements in excess returns associated with past movements in

$\widehat{cay}_t$ may provide a better proxy for movements in the conditional first moment of excess returns than that available using other popular forecasting variables. Second, it is important for theoretical modeling of asset pricing behavior to know whether the same variables that forecast returns also forecast variances and at the same frequencies, revealing whether the first and second moments of returns are in fact related. Third, because $\widehat{cay}_t$ relies on information on aggregate consumption and labor earnings, its use as a predictive variable presents a new opportunity to investigate the empirical linkages between the real economy and financial market volatility. We do so here, using quarterly and monthly data on $\widehat{cay}_t$.

To obtain a measure of volatility for the excess return on the CRSP-VW index, we use the time-series variation of daily returns:

$$V_t = \widehat{\sigma}_{st} = \sqrt{\sum_{k \in t} (R_{sk} - \overline{R}_s)^2},$$

where V_t is the sample volatility of the market return in period t, R_{sk} is the daily CRSP-VW return minus the implied daily yield on the 3 month Treasury bill rate, $\overline{R}_s$ is the mean of R_{sk} over the whole sample, k represents a day, and t is either a quarter or a month. Next we forecast R_{st+1} and V_{t+1} with the log consumption-wealth ratio proxy, $\widehat{cay}_t$ and a range of other conditioning variables used in the existing literature, and store the fitted values to obtain a measure of the conditional mean of returns, $E_t(R_{st+1})$, the conditional volatility of returns, $E_t V_{t+1}$, respectively. Our estimate of the conditional Sharpe ratio is simply $\frac{E_t(R_{st+1})}{E_t V_{t+1}}$. For the CRSP-VW index, the quarterly mean excess return is 0.019; the quarterly standard deviation is 0.0817.

Table 4 presents long-horizon regressions of volatility, V_{t+h}, for several horizons, h, on a variety of predictive variables. The first six rows report results using quarterly data; the last two rows using monthly data. There is substantial autocorrelation in measured volatility, thus we include two lags of volatility in our forecasting equations for V_t; the results of estimating a purely autoregressive specification are reported in row 1. Past volatility is a statistically significant predictor of future volatility up to four quarters ahead, with adjusted R^2 statistics monotonically declining from 22 percent at a one quarter horizon. At a horizon of six quarters, past volatility explains virtually nothing of future volatility.

The second and third rows of Table 4 display the forecasting power of the consumption-wealth ratio proxy, $\widehat{cay}_t$, for future volatility using quarterly data. Two aspects of these findings stand out. First, the signs of the significant coefficients in these regressions are all negative. Recalling that high values of $\widehat{cay}_t$ predict high excess returns (Table 1), this result implies that conditional expected excess returns are *negatively* related to conditional volatility. The finding suggests that stock market volatility by itself is a poor proxy for variation in the equity risk premium, since high risk-premia cannot be explained by high stock market volatilty and vice versa. Second, the regression results indicate that $\widehat{cay}_t$ is both a statistically significant and economically important determinant of future stock market volatility. When $\widehat{cay}_t$ is the sole predictive variable (row 2), it is statistically significant at the 5 percent level, over horizons ranging from one to 12 quarters, with R^2 statistics starting at 12 percent for a one quarter horizon and rising to a peak of 19 percent six quarters ahead. The marginal predictive power of $\widehat{cay}_t$ survives when past volatility is controlled for (row 3), and including in the purely autoregressive specification allows the regression to explain an additional 16 percent of the variation in volatility six quarters ahead (compare row 3 to row

1).

As an illustration of the extent to which these results imply a negative correlation between the conditional first and second moments of excess returns, we compute the correlation between the fitted values from the regression displayed in row 3 of Table 4 and the fitted values from a regression of excess returns on lagged $\widehat{cay}_t$. This correlation is found to be -0.956.[11] Thus, the conditional expected excess stock return is strongly negatively correlated with conditional volatility, according to these estimates. These results also demonstrate that volatility is predictable by the at least some of the same variables that predict excess returns, contrary to common perception that this is not the case (e.g., Cochrane (2001)).

Our finding that conditional expected excess returns are strongly negatively related to conditional volatility is consistent with recent empirical evidence by Brandt and Kang (2001). Brandt and Kang take an approach that is quite different from that employed in this chapter. Instead of forecasting returns and volatility with particular conditioning variables, they model the conditional mean and conditional volatility of stock returns as latent variables which follow a bivariate Gaussian first-order VAR process. Thus, rather than assuming that the dynamics of conditional moments are determined by specific conditioning variables, they make assumptions about the joint time-series process of the unobservable conditional mean and conditional variance, and use a filtering algorithm to estimate a VAR for these variables, given an assumed distribution of the VAR innovations. Despite the divergence in approach from that take here, the result is the same. Brandt and Kang find that the correlation

[11] Although this correlation may seem quite striking, it is less surprising when one considers that the unconditional correlation between the excess return and volatility of the CRSP-VW return is also negative, equal to -0.356.

between the VAR innovations of the mean and volatility, which condition on past values of the latent variables, is negative and statistically significant, with a correlation of -0.63 using monthly data. Interestingly, however, Brandt and Kang also find that the *unconditional* correlation between the latent first and second moments is positive. They argue that much of the disagreement in the literature over the sign of the conditional correlation might be explained by the possibility that some studies may have inadvertently measured something closer to the unconditional correlation, either because the conditioning variables or empirical specifications of volatility made inadequate allowances for time-variation in the conditional moments of returns.

The fourth row of Table 4 uses the dividend-price ratio to forecast volatility. The coefficient on this variable, like that on $\widehat{cay}_t$, is negative, and it is statistically significant up to six quarters ahead, again suggesting a negative correlation between expected returns volatility. But row 5 of Table 4 shows that the predictive power of the dividend-yield is driven out by $\widehat{cay}_t$. In addition, other results (not reported) indicated that the predictive power of $d_t - p_t$ for future volatility is quite sensitive to the sample used. In particular, eliminating just the last two years of data renders the estimated coefficients on the dividend-yield statistically insignificant at conventional significant levels.[12] We did not find such instability using $\widehat{cay}_t$ as a predictive variable for volatility.

The sixth row adds three additional regressors to the set of forecasting variables for volatility: DEF_t, a commercial paper-Treasury spread, CP_t, and the one year Treasury yield, $TB1Y_t$. The last three predictive variables are those used by Whitelaw (1994) to

[12] The insignificant coefficients found when data is used up through 1999 are consistent with the findings of other researchers investigating the link between $d_t - p_t$ and future volatility (e.g., Campbell (2001)).

forecast volatility at monthly and quarterly horizons. Although these variables are not strong predictors of excess returns, it is nevertheless worth checking whether the forecasting power of $\widehat{cay}_t$ for future volatility is robust to the inclusion of these additional regressors. Row 6 shows that, in this multivariate regression, all variables except the default spread have marginal predictive power at one horizon or another, with $\widehat{cay}_t$, $d_t - p_t$, and CP_t displaying forecasting power at horizons less than six quarters, and $TB1Y_t$ displaying forecasting power at horizons in excess of three years.

Rows 7 and 8 of Table 4 present results from forecasting regressions for V_t using monthly data; row 7 presents the pure autoregressive regression; row 8 includes $\widehat{cay}_t$ as a predictive variable. The results are qualitatively similar to those using quarterly data, with $\widehat{cay}_t$ adding most to the predictive capacity of the regression at horizons in excess of 12 months.

Several aspects of the results in Table 4 are worth emphasizing. First, the findings imply that the conditional Sharpe ratio cannot be constant since variation in the conditional mean, in the numerator, moves in the opposite direction of variation in conditional volatility, in the denominator. Second, the negative correlation between the conditional mean and conditional volatility documented here is inconsistent with leading equilibrium asset pricing models that are capable of generating a countercyclical price of risk (e.g., Campbell and Cochrane (1999); Barberis, Huang, and Santos (2001)). Rather than a negative correlation, these models predict a positive correlation between the conditional mean and conditional volatility, and they generate a countercyclical Sharpe ratio only because there is more variation in the mean than in the volatility. One theoretical framework that *can* generate a negative correlation between the conditional first and second moments of returns is the the model considered in Whitelaw (2000). Whitelaw assumes that consumption growth follows a Markov regime switching process with time-varying transitory probabilities and shows that such a structure can generate a negative correlation between stock market volatility and expected returns. An important difficulty with this Markov regime-switching framework, however, is that it does not deliver persistent price-dividend ratios, nor does it generate long-horizon forecastability of excess returns by the consumption-wealth ratio, as documented in Table 1.

A third noteworthy aspect of the results in Table 4 is the mere finding that conditional volatility varies: while there is a vast literature documenting time-variation in stock market volatility at high frequencies, it is often thought that volatility is not strongly forecastable at frequencies as low as a quarter (e.g., Christoffersen and Diebold (2000); Campbell (2001)). Table 4 demonstrates that this is not the case. Instead, volatility is strongly forecastable by $\widehat{cay}_t$, at horizons ranging from one quarter to three years. Finally, we note that these results on the time-series variation in volatility can be linked to the literature on cross-sectional variation in stock returns. For example, Lettau and Ludvigson (2001b) investigate a conditional version of a consumption-based capital asset pricing model (CAPM), using $\widehat{cay}_t$ as a conditioning variable. The argument for using conditioning information is that the consumption beta in this model should depend on the conditional Sharpe ratio for the market portfolio. The use of $\widehat{cay}_t$ as a conditioning variable was motivated by the finding in Lettau and Ludvigson (2001a) that $\widehat{cay}_t$ captures time-variation in expected excess returns, suggesting that it may also proxy for movements in the Sharpe ratio as long as volatility does not move one-for-one with expected returns. The results presented in this chapter confirm that volatility does not move one-for-one with expected returns and therefore bolster the case for using $\widehat{cay}_t$ as a conditioning variable in cross-sectional asset pricing tests where time-variation in the price of risk is important.

Figure 1 plots our estimate over time of the conditional volatility of the excess return to the CRSP-VW index. The figure plots the fitted values from the regression specification given in row 3 of Table 4, which includes $\widehat{cay}_t$ and two lags of volatility as predictors of quarterly volatility. NBER dated recessions are indicated with shaded bars. Consistent with findings in Brandt and Kang (2001), the figure suggests that conditional volatility is high in recessions, but falls through the course of the recession when expected returns are rising. By contrast, conditional volatility tends to rise over the course of an expansion when conditional expected returns are falling, illustrating the negative correlation between expected returns and conditional volatility that is evident in the regression coefficients reported in Tables 1 and 4.

4 The Conditional Sharpe Ratio

It is well known that conditional expected returns are countercyclical. The finding reported above, namely that the conditional mean excess return is negatively related to conditional volatility, suggests that there will be pronounced countercyclical variation in the Sharpe ratio. The estimated value of this ratio over time is plotted in Figure 2 for quarterly excess returns on the CRSP-VW stock index.[13] The figure confirms that the Sharpe ratio, plotted on a quarterly basis, is countercyclical, falling over the course of an expansion and shooting

[13]We forecast $R_{st+1} - R_{ft+1}$ using the log consumption-wealth ratio proxy, $\widehat{cay}_t$; we forecast V_t using cay_t and two lags of v_t.

up at the beginning of recessions. Note that there are also a few periods during which the conditional Sharpe ratio is estimated to be negative. This occurs because our estimate of conditional expected returns–the fitted values from a regression of excess returns on lagged $\widehat{cay}_t$–are occasionally negative. This result stems from the linear regression specification underlying our identification of expected returns, and is not unique to the use of any particular forecasting variable. Nevertheless, its worth noting that an occasional negative risk premium on stock market wealth is not necessarily inconsistent with equilibrium asset pricing models in which the covariance of consumption growth with the stochastic discount factor varies over time (Boudoukh, Richardson, and Whitelaw (1997); Whitelaw (2000)).

The magnitude of countercyclical variability in the Sharpe ratio displayed in Figure 2 is not well captured by leading equilibrium asset pricing models. As an illustration of the existing theoretical gap, Figure 2 also plots the implied Sharpe ratio from one of the leading paradigms for rationalizing observed asset pricing behavior: the model explored in Campbell and Cochrane (1999). The Campbell-Cochrane model is a habit persistence framework in which utility takes the form $u(C_t, X_t) = \frac{(C_t - X_t)^{1-\gamma}}{1-\gamma}$, where X_t is the external consumption habit. The Sharpe ratio predicted by the Campbell-Cochrane model, which we denote SR_t^{CC}, is a nonlinear function of consumption growth, and takes the form

$$SR_t^{CC} = \{e^{\gamma^2 \sigma^2 [1+\lambda(s_t)]^2} - 1\}^{1/2} \approx \gamma\sigma[1 + \lambda(s_t)], \tag{4}$$

where, s_t is the log of the surplus consumption ratio, defined $S_t \equiv \frac{C_t - X_t}{C_t}$, and $\lambda(s_t)$ is the sensitivity function specified in Campbell and Cochrane. The log surplus consumption ratio evolves as a heteroskedastic, first-order autoregressive process:

$$s_{t+1} = (1-\phi)\overline{s} + \phi s_t + \lambda(s_t)(\Delta c_{t+1} - g), \tag{5}$$

where g is the mean rate of consumption growth and ϕ is the persistence of the habit stock. It is straightforward to compute the implied Sharpe ratio of the Campbell-Cochrane model by combining (4) and (5) with data on aggregate consumption .[14] This series is plotted in Figure 2 along side our estimate of the Sharpe ratio over time.

Although Campbell and Cochrane (1999) show that the model they study does a reasonable job of matching variation in the first moment of excess returns, Figure 2 suggests that the model produces an unrealistically small amount of countercylical variation in the Sharpe ratio. This occurs, in part, because conditional volatility (in the denominator) counterfactually moves in the same direction as the conditional mean (in the numerator). For example, the estimated Sharpe ratio for excess returns on the CRSP-VW index, SR_t^{VW}, ranges from -0.45 to 1.76 on a quarterly basis. By contrast, SR_t^{CC} ranges from 0 and 0.4.

Other equilibrium models also fail to produce the observed degree of variability in SR_t^{VW}. Barberis, Huang, and Santos (2001) study an economy in which investors derive utility from consumption and wealth, and show that this model can replicate persistent time-variation in conditional excess returns. Like the Campbell-Cochrane model, however, the Sharpe ratio they report ranges from about 0.20 to 0.40 on a quarterly basis, far less than that documented in Figure 2. Even the Whitelaw (2000) model, which in principle can rationalize a negative correlation between the conditional mean and conditional volatility of stock re-

[14] We use the value of $\overline{s}$ calibrated in Campbell and Cochrane (1999).

turns, generates only a very small fraction of the observed variation in the Sharpe ratio, with quarterly values ranging from a low of -0.0012 to a high of 0.0122.[15]

The shortcomings of existing equilibrium models documented here are distinct from those underlying the "equity premium puzzle" of Mehra and Prescott (1985) and Hansen and Jagannathan (1991). Those studies show that standard asset pricing theory fails to account for the high mean value of the Sharpe ratio. While those papers focused on the average value of the Sharpe ratio, we concentrate here on its variation through time. The evidence presented in this chapter suggests that even our best fitting asset pricing models have difficulty replicating the observed pattern of variation in the price of stock market risk, and leave a "Sharpe ratio volatility puzzle" that remains to be explained.

The theoretical models discussed above are all consumption-based asset pricing frameworks that generate movements in the Sharpe ratio from movements in either risk-aversion (i.e., Campbell and Cochrane (1999); Barberis, Huang, and Santos (2001)), or from movements in the conditional correlation between stock returns and the intertemporal marginal rate of substitution in consumption (i.e., Whitelaw (2000)). Nevertheless, there is another possible channel through which variability in the Sharpe ratio can be generated in consumption-based models: time-variation in conditional volatility of consumption growth. For example, consider a popular time-separable specification in which investors have constant relative risk aversion utility taking the form $u(C_t) = \frac{(C_t)^{1-\gamma}}{1-\gamma}$; and where, in the limit as $\gamma \to 1$, $u(c_t) = \log(c_t)$. In this case, the investor's first-order condition for optimal consumption choice is an Euler equation relating excess stock returns to the marginal rate of

[15]These statements are based on the numbers reported in Figure 6 of Barberis, Huang, and Santos (2001) and Table 3 of Whitelaw (2000).

substitution in consumption:

$$1 = \beta E_t \left(R_{st+1} \frac{c_{t+1}^{-\gamma}}{c_t^{-\gamma}} \right), \tag{6}$$

where β is the subjective rate of time-preference, R_{st+1} is the net return on stocks, and $m_{t+1} \equiv \beta \frac{c_{t+1}^{-\gamma}}{c_t^{-\gamma}}$ is the marginal rate of substitution in consumption. Applying a covariance decomposition to (6), and using $R_{ft} = 1/E_t(m_{t+1})$, the risk-premium on stocks is given by $E_t(R_{st+1}) - R_{ft} = -R_{ft}\text{cov}_t(m_{t+1,}R_{st+1})$. If we further assume that consumption growth is lognormally distributed, we obtain the following approximate expression for the conditional Sharpe ratio in this model:

$$\frac{E_t R_{st+1} - R_{ft}}{\sigma_t(R_{st+1})} \approx \gamma \sigma_t(\Delta c_{t+1}) \rho_t(m_{t+1}, R_{st+1}). \tag{7}$$

The left-hand-side of (7) is the conditional Sharpe ratio for excess stock returns. The numerator is the conditional expected excess return; the denominator is the conditional standard deviation of stock returns. This expression says that the conditional Sharpe ratio is equal to γ, the coefficient of relative risk aversion, times $\sigma_t(\Delta c_{t+1})$, the conditional standard deviation of consumption growth, times $\rho_t(m_{t+1}, R_{st+1})$, the conditional correlation between the intertemporal marginal rate of substitution, m_t, and the return on stocks.

Equation (7) highlights the mechanisms by which the models discussed above generate time-variation in the conditional Sharpe ratio. In Campbell and Cochrane (1999) and Barberis, Huang, and Santos (2001), the Sharpe ratio moves over time because risk aversion, γ, is not constant but time-varying. In Whitelaw (2000), time-variation in the Sharpe ratio is generated by time-variation in the conditional correlation, $\rho_t(m_{t+1}, R_{st+1})$. As the approximate framework in (7) makes clear, however, time-variation in the risk-return tradeoff could also be generated by variability in $\sigma_t(\Delta c_{t+1})$, even if risk aversion or the conditional correlation, $\rho_t(m_{t+1}, R_{st+1})$, are constant.

Can the magnitude of variation in the Sharpe ratio be explained by time-variation in consumption volatility? Within the confines of this chapter it is not possible to investigate the range of possible econometric techniques for modeling changing volatility in consumption growth. Instead, we make a first-pass at addressing this question by modeling the volatility of consumption growth as a generalized autoregressive conditional heteroskedasticity (GARCH) process. With these estimates of of $\sigma_t(\Delta c_{t+1})$ in hand, we then move on to ask whether the framework in (7) is helpful in explaining the pattern of variability in the Sharpe ratio that we document here.

We assume that the conditional correlation, $\rho_t(m_{t+1}, R_{st+1})$, is 1, and choose risk aversion, γ, to match the mean Sharpe ratio relative to the mean of our estimate of $\sigma_t(\Delta c_{t+1})$. Notice that any portfolio that is sufficiently diversified (a mean-variance efficient portfolio) will have $\rho_t(m_{t+1}, R_{st+1}) = 1$. Although a broad stock market return may not be an efficient portfolio, setting this correlation to one provides a reasonable benchmark because many asset pricing studies implicitly assume that such an asset is a highly correlated with an efficient portfolio and set this correlation to 1 in undertaking calibration exercises.[16] In addition, this approach allows us to isolate the contribution of consumption risk in explaining the pattern of variability in the risk-return tradeoff. Thus, the Sharpe ratio we measure for the consumption

[16] Furthermore, Campbell (2001) and Cochrane (2001) emphasize that this correlation is hard to measure accurately because estimates are sensitive to data definition, measurement error, the length of the horizon, and data aggregation.

model in (7) is simply $\gamma\, \widehat{\sigma}_t(\Delta c_{t+1})$, where $\widehat{\sigma}_t(\Delta c_{t+1})$ denotes the estimated volatility measure from the GARCH procedure. We refer to this framework as the *consumption volatility* model and denote the Sharpe ratio implied by this model as $SR_t^{CV} \equiv \gamma\, \widehat{\sigma}_t(\Delta c_{t+1})$.

Table 5 presents maximum likelihood estimates of a GARCH(1,1) process for the volatility of the innovation of quarterly consumption growth. The GARCH model takes the form

$$
\begin{aligned}
\Delta c_t &= \alpha_0 + \sum_{i=1}^{3} \alpha_i \Delta c_{t-1} + \epsilon_t \\
\sigma_t^2 &= \delta_0 + \delta_1 \epsilon_t^2 + \delta_2 \sigma_{t-1}^2 + \delta_3 \mathbf{X}_{t-1},
\end{aligned}
$$

where σ_t^2 is the conditional variance of ϵ_t, and $\mathbf{X}_{t-1}$ is a vector of predetermined conditioning variables that may influence the volatility of consumption growth. The table reports estimates of the parameters α_i and δ_i for four specifications: one with no conditioning variables (column 1); one in which $\mathbf{X}_{t-1} = \widehat{cay}_{t-1}$ (column 2); one in which $\mathbf{X}_{t-1} = r_{st-1} - r_{ft-2}$ (column 3); and one in which $\mathbf{X}_{t-1} = (\widehat{cay}_{t-1},\, r_{st-1} - r_{ft-2})'$ (column 4). The results suggests that the volatility of consumption growth is not constant over time; for example the coefficient on the GARCH term , δ_2, is much larger than its standard deviation. Moreover, columns 2-4 indicate that both $\widehat{cay}_{t-1}$and $r_{st-1} - r_{ft-2}$ have marginally significant explanatory power for consumption volatility when they are included as regressors either by themselves or together. Thus we take the fitted values of σ_t^2 from the fourth column, in which $\mathbf{X}_{t-1} = (\widehat{cay}_{t-1},\, r_{st-1} - r_{ft-2})'$, as our estimate of the conditional variance of consumption growth. The square root of these fitted values, $\sqrt{\widehat{\sigma}_t^2}$, is our estimate of conditional volatility, used to compute SR_t^{CV}.

The value of relative risk aversion, γ, that matches the mean Sharpe ratio in our sample is is 92, a large number that illustrates the equity premium puzzle emphasized by Mehra and Prescott (1985) and Hansen and Jagannathan (1991).[17] The focus of this paper is not on this unconditional puzzle, but on the pattern of variability in the conditional Sharpe ratio. Nevertheless the high value for risk aversion required to match the mean Sharpe ratio underscores an important point, namely that modeling the variance of consumption growth as time-varying does not by itself help resolve the equity premium puzzle. Although the results in Table 5 imply that there may be some variation in the volatility of consumption growth, it is quantitatively minuscule when compared to the variability of SR_t^{VW}.

Table 6 presents summary statistics for the empirical Sharpe ratio estimated from the data, SR_t^{VW}, the Campbell-Cochrane Sharpe ratio, SR_t^{CC}, and the consumption-volatility Sharpe ratio, SR_t^{CV}. The table illustrates several important aspects of the Sharpe ratio volatility puzzle. First, the standard deviation of SR_t^{VW} is over five times as large as that of either SR_t^{CC} or SR_t^{CV}, reinforcing the notion that consumption-based models fail to replicate the magnitude of volatility in the risk-return tradeoff. Second, the Campbell-Cochrane Sharpe ratio is too autocorrelated, whereas the consumption-volatilty model produces about the right autocorrelation. Third, SR_t^{CC} is positively correlated with SR_t^{VW} with this correlation equal to about 0.4 in the data. By contrast, the consumption volatility model fails miserably along this dimension, displaying a *negative* correlation, equal to -0.3 with SR_t^{VW}.

[17]The mean Sharpe ratio in our sample is 0.78 on an annual basis, somewhat larger than that typically reported (for example, Campbell and Cochrane (1999) report a Sharpe ratio for log returns of 0.43 in post-war data). As a results, the value for γ needed to match this Sharpe ratio is also somewhat larger than that typically required of the consumption-based model considered above. The reason is that we compute volatility, in the denominator, from daily returns and then convert to a quarterly rate. Because daily returns are positively serially correlated, this number is smaller than the volatility of quarterly or monthly returns. Computing volatility from either of the latter delivers a value for γ that is closer to 50, rather than the 92 we report above.

This negative correlation is evident in Figure 3, which plots SR_t^{VW} and SR_t^{CV} over time. In short, time-variation in consumption volatility appears unhelpful in explaining observed variability in the risk-return tradeoff on broad stock returns. And, although we have not explored other methodologies for estimating time-varying consumption risk, (for example, stochastic volatility), it seems unlikely that such an extensions would produce significantly more volatile consumption growth, or volatility that varied over time in a drastically different manner. We therefore conclude that variation in consumption risk alone is unlikely to help resolve the Sharpe ratio volatility puzzle.

In summary, the evidence reported here presents a significant challenge to existing asset pricing theory. Of the leading equilibrium paradigms capable of replicating any time series variation in conditional moments, most predict that the conditional mean and conditional volatility of excess returns move in the same direction, inconsistent with our own findings and those in a number of other recent papers. Even those models that can rationalize such a negative correlation (e.g., Whitelaw (2000)) do not solve the larger puzzle, namely that all of these models completely miss the sheer magnitude of volatility in the risk-return tradeoff. At the same time, the results merely serve to strengthen some key conclusions of the existing literature, namely that predictability of excess returns cannot be explained by movements in stock market volatility, but must instead be explained by movements in risk-aversion, or in the conditional covariance of consumption growth and returns, or some combination of the two.

5 Conclusion

There is now a large and growing body of empirical evidence that finds forecastability of excess equity returns and measures of their volatility. Recent theoretical work in financial economics has demonstrated that such forecastability is not necessarily inconsistent with market efficiency. In particular, stock market predictability can be generated by time-variation in the rate at which rational, utility maximizing investors discount expected future cash-flows from risky assets. These theoretical advances hold out hope that a unified framework for rationalizing variation in the risk-return tradeoff can be developed.

This chapter reviews what is known about the time-series variability in the expected excess return on the stock market, relative to its conditional volatility. We examine the empirical procedures and results of a large number of studies that canvass the subject of predictability in stock returns and stock return volatility, and we assess whether the current state of theoretical knowledge can account for such predictability. We also present some new empirical evidence using a proxy for the log consumption-aggregate wealth ratio as a forecaster of both the mean and volatility of excess stock returns.

Although the existing assemblage of empirical work has spawned some measure of disagreement about the time-series properties of broad stock market returns, our own assessment of this literature is that the balance of empirical evidence favors significant long-horizon predictability of excess stock returns, and a negative correlation between the conditional mean and conditional volatility of these returns. Indeed, our own investigation implies that both excess stock returns and the volatility of returns are forecastable by a proxy for the log consumption-wealth ratio, and that the conditional mean is strongly negatively correlated with the conditional standard deviation.

This evidence presents a considerable challenge for existing asset pricing theory. Leading asset pricing models–those capable of generating any variation in the risk-return tradeoff–typically do not imply that times of high equity risk premia coincide with times of low stock market volatility, and vice versa. More significantly, these models leave a "Sharpe ratio volatility puzzle"that remains to be explained: not only is the Sharpe ratio high on average as is well understood, it is characterized by pronounced countercyclical variation that is not matched in magnitude by existing asset pricing models.

Appendix A: Data Description

CONSUMPTION, C_t

Consumption is measured as expenditure on nondurables and services, excluding shoes and clothing. The quarterly data are seasonally adjusted at annual rates, in billions of chain-weighted 1996 dollars. The components are chain-weighted together, and this series is scaled up so that the sample mean matches the sample mean of total personal consumption expenditures. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

AFTER-TAX LABOR INCOME, Y_t

Labor income is defined as wages and salaries + transfer payments + other labor income - personal contributions for social insurance - taxes. Taxes are defined as [wages and salaries/(wages and salaries + proprietors' income with IVA and Ccadj + rental income + personal dividends + personal interest income)] times personal tax and nontax payments, where IVA is inventory valuation and Ccadj is capital consumption adjustments. The quarterly data are in current dollars. A real per capita series is created by dividing by a measure of the population and the price deflator listed below. Our source is the Bureau of Economic Analysis.

POPULATION

A measure of population is created by dividing real total disposable income by real per capita disposable income. Our source is the Bureau of Economic Analysis.

WEALTH, A_t

Total wealth is household net wealth in billions of current dollars, measured at the end of the period. We lag this series one period to produce a measure of beginning of period wealth. Stock market wealth includes direct household holdings, mutual fund holdings, holdings of private and public pension plans, personal trusts, and insurance companies. Nonstock wealth is the residual of total wealth minus stock market wealth, and includes ownership of privately traded companies in noncorporate equity. Our source is the Board of Governors of the Federal Reserve System. A real per capita series is created by dividing by a measure of the population and the price deflator listed below.

PRICE DEFLATOR

The nominal after-tax labor income and wealth data are deflated by the personal consumption expenditure chain-type deflator (1996=100), seasonally adjusted. Our source is the Bureau of Economic Analysis.

EXCESS RETURNS, $r_{t+1} - r_{ft}$

Excess returns are returns to the CRSP value-weighted stock index, less the 3-month treasury bill yield. Our sources are the Center for Research in Securities Prices and the Board of Governors of the Federal Reserve System.

CRSP DIVIDEND-PRICE RATIO, $d_t - p_t$

The CRSP Dividend-Ratio is calculated as the log ratio of CRSP dividends to the price level of the CRSP value-weighted stock index (imputed from CRSP-VW returns, including dividends). Our source is the Center for Research in Securities Prices.

DEFAULT SPREAD, DEF_t

The default spread is the difference between the BAA corporate bond rate and the AAA corporate bond rate. Our source is the Moody's Corporate Bond Indices.

RELATIVE BILL RATE, $RREL_t$

The relative bill rate is the 3-month treasury bill yield less its four-quarter moving average. Our source is the Board of Governors of the Federal Reserve System.

TERM SPREAD, TRM_t

The term spread is the difference between the 10-year treasury bond yield and the 3-month treasury bill yield. Our source is the Board of Governors of the Federal Reserve System.

COMMERCIAL PAPER SPREAD, CP_t

The commercial paper spread is the difference between the yield on 6-month commercial paper and the 3-month treasury bill yield. Our source is the Board of Governors of the Federal Reserve System.

ONE-YEAR TREASURY BILL YIELD, $TB1Y_t$

Our source for the 1-year treasury bill yield is the Board of Governors of the Federal Reserve System.

References

Ang, A., and G. Bekaert (2001): "Stock Return Predictability: Is it There?," NBER Working Paper No. 8207.

Barberis, N., M. Huang, and T. Santos (2001): "Prospect Theory and Asset Prices," *Quarterly Journal of Economics*, 116(1), 1–53.

Bollerslev, T., R. Engle, and J. Wooldridge (1988): "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy*, 96, 116–131.

Bossaerts, P., and P. Hillion (1999): "Implementing Statistical Criteria to Select Return Forecasting Models: What Do We Learn?," *Review of Financial Studies*, 12, 405:428.

Boudoukh, J., M. Richardson, and R. F. Whitelaw (1997): "Nonlinearities in the Relation between the Equity Risk Premium and the Term Structure," *Management Science*, 43(3), 371–385.

Brandt, M. W., and Q. Kang (2001): "On the Relation Between the Conditional Mean and Volatility of Stock Returns: A Latent VAR Approach," Unpublished paper, University of Pennsylvania.

Breen, W., L. R. Glosten, and R. Jagannathan (1989): "Economic Significance of Predictable Variations in Stock Index Returns," *Journal of Finance*, 44(5), 1177–89.

Campbell, J. Y. (1987): "Stock Returns and the Term Structure," *Journal of Financial Economics*, 18, 373–399.

——— (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

——— (2001): "Consumption-Based Asset Pricing," Unpublished paper, Harvard University.

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

Campbell, J. Y., A. W. Lo, and C. MacKinlay (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.

Campbell, J. Y., and R. J. Shiller (1988): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195–227.

——— (2001): "Valuation Ratios and the Long-Run Stock Market Outlook: An Update," NBER working paper No. 8221.

Christoffersen, P. F., and F. X. Diebold (2000): "How Relevant is Volatility Forecasting for Financial Risk Management?," *The Review of Economics and Statistics*, 82, 12–23.

Clark, T., and M. McCracken (1999): "Tests of Equal Forecast Accuracy and Encompassing for Nested Models," Unpublished Paper, Federal Reserve Bank of Kansas City.

Cochrane, J. H. (1991a): "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, 5(2), 243–280.

——— (1991b): "Volatility Tests and Efficient Markets: A Review Essay," *Journal of Monetary Economics*, 127, 463–485.

——— (1994): "Permanent and Transitory Components of GDP and Stock Prices," *Quarterly Journal of Economics*, 109, 241–265.

——— (1997): "Where is the Market Going? Uncertain Facts and Novel Theories," *Economic Perspectives, Federal Reserve Bank of Chicago*, 11, 1–37.

——— (2001): *Asset Pricing*. Princeton University Press, Princeton, NJ.

Constantinides, G., and D. Duffie (1996): "Asset Pricing With Heterogeneous Consumers," *Journal of Political Economy*, 104, 219–40.

Fama, E. F., and K. R. French (1988): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–27.

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

Ferson, W., S. Sarkissian, and T. Simin (1999): "Spurious regressions in financial economics?," Unpublished paper, University of Washington.

French, K., G. W. Schwert, and R. F. Stambaugh (1987): "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3–30.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993): "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 48(5), 1779–1801.

Goyal, A., and I. Welch (1999): "Predicting the Equity Premium," Unpublished paper, University of California Los Angeles.

Hansen, L. P., and R. Jagannathan (1991): "Restrictions on Intertemporal Marginal Rates of Substitution Implied by Asset Returns," *Journal of Political Economy*, 99, 225–262.

Harvey, C. R. (1989): "Time-Varying Conditional Covariances in Tests of Asset Pricing Models," *Journal of Financial Economics*, 24, 289–371.

Harvey, D., S. Leybourne, and P. Newbold (1998): "Tests for Forecast Encompassing," *Journal of Business and Economic Statistics*, 16, 254–259.

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.

Kandel, S., and R. Stambaugh (1989): "Modeling Expected Stock Returns for Long and Short Horizons," Unpublished Paper , Rodney L. White Center, Wharton School, University of Pennsylvania.

Lamont, O. (1998): "Earnings and Expected Returns," *Journal of Finance*, 53, 1563–87.

LeRoy, S., and R. Porter (1981): "The Present Value Relation: Tests Based on Variance Bounds," *Econometrica*, 49, 555–557.

LETTAU, M., AND S. C. LUDVIGSON (2001a): "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56(3), 815–849.

——— (2001b): "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia are Time-Varying," *Journal of Political Economy,* forthcoming.

——— (2001c): "Time-Varying Risk Premia and the Cost of Capital: An Alternative Implication of the Q Theory of Investment," *Journal of Monetary Economics,* forthcoming.

——— (2001d): "Understanding Trend and Cycle in Asset Values: Bulls, Bears, and the Wealth Effect on Consumption," Unpublished Paper, New York University.

LEWELLEN, J. W. (1999): "The Time Series Relations Among Expected Return, Risk, and book-to-market," *Journal of Financial Economics*, 54, 5–53.

LINTNER, J. (1965): "Security Prices, Risk and Maximal Gains from Diversification," *Journal of Finance*, 20, 587–615.

LUDVIGSON, S., AND C. STEINDEL (1999): "How Important Is the Stock Market Effect on Consumption?," *Economic Policy Review, Federal Reserve Bank of New York*, 5, 29–51.

MANKIW, N. G. (1982): "Hall's Consumption Hypothesis and Durable Goods," *Journal of Monetary Economics*, 10(3), 417–425.

MCCRACKEN, M. (1999): "Asymptotics for Out-of-Sample Tests of Causality," Unpublished manuscript, Louisiana State University.

MEHRA, R., AND E. PRESCOTT (1985): "The Equity Premium Puzzle," *Journal of Monetary Economics*, 15, 145–161.

NELSON, C. C., AND M. J. KIM (1993): "Predictable Stock Returns: The Role of Small Sample Bias," *Journal of Finance*, 43, 641–661.

NEWEY, W. K., AND K. D. WEST (1987): "A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

RICHARDSON, M., AND J. STOCK (1989): "Drawing Inferences from Statistics Based on Multi-Year Asset Returns," *Journal of Financial Economics*, 25, 323–348.

SCHWERT, G. W. (1989): "Why Does Stock Market Volatility Change over Time," *Journal of Finance*, 44(5), 1115–53.

SHARPE, W. (1964): "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance*, 19, 425–444.

SHILLER, R. J. (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 71, 421–436.

SHLEIFER, A. (1995): "Comments on Stock Ownership Patterns, Stock Market Fluctuations, and Consumption," *Brookings Papers on Economic Activity*, (2), 295–372.

SIEGEL, J. J. (1998): *Stocks for the Long Run: The Definitive Guide to Financial Market Returns and Long-term Investment Strategies*. McGraw Hill, New York, NY, 2 edn.

STAMBAUGH, R. F. (1999): "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421.

STOCK, J. H., AND M. WATSON (1993): "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 61, 783–820.

VALKANOV, R. (2001): "Long-Horizon Regressions: Theoretical Results and Applications," *Journal of Financial Economics,* forthcoming.

VUOLTEENAHO, T. (2000): "Understanding the Aggregate Book-Market Ratio and its Implications to Current Equity-Premium Expectations," Unpublished manuscript, Harvard University.

WHITELAW, R. F. (1994): "Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns," *The Journal of Finance*, 49(2), 515–41.

——— (2000): "Stock Market Risk and Return: An Equilibrium Approach," *Review of Financial Studies*, 13(3), 521–547.

Table 1

Forecasting Quarterly Stock Returns, 1952:4-2000:4

#	Constant (t -stat)	lag (t -stat)	$\widehat{cay}_t$ (t -stat)	$d_t - p_t$ (t -stat)	$RREL_t$ (t -stat)	TRM_t (t -stat)	DEF_t	$\overline{R}^2$
Panel A: Real Returns; 1952:4-2000:4								
1	**0.03**	0.04						-0.00
	(4.90)	(0.57)						
2	**0.04**		**1.78**					0.06
	(7.30)		(2.95)					
3	**0.05**	-0.05	**1.88**					0.06
	(7.07)	(-0.80)	(3.09)					
Panel B: Excess Returns; 1952:4-2000:4								
4	**0.02**	0.07						-0.00
	(3.03)	(1.02)						
5	**0.03**		**1.80**					0.07
	(5.38)		(3.42)					
6	**0.04**	-0.01	**1.83**					0.07
	(5.27)	(-0.21)	(3.36)					
Panel C: Additional Controls; Excess Returns; 1952:4-2000:4								
7	0.11			0.03				0.01
	(1.60)			(1.33)				
8	0.02		**1.86**	-0.00				0.07
	(0.27)		(3.55)	(-0.20)				
9	0.00	-0.08	**1.89**	-0.01	**-2.50**	-0.36	0.00	0.09
	(0.03)	(-1.38)	(3.54)	(-0.40)	(-2.98)	(-0.47)	(0.21)	

Notes: See next page.

Notes to Table 1

The table reports estimates from OLS regressions of stock returns on lagged variables named at the head of a column. The dependent variable is the log of the return on the CRSP Value-Weighted stock market index. The regressors are as follows: lag denotes a one-period lag of the dependent variable, $\widehat{cay}_t \equiv c_t - \widehat{\beta}_a a_t - \widehat{\beta}_y y_t$, where c_t is consumption, a_t is asset wealth and y_t is labor income, $d_t - p_t$ is the log dividend-price ratio; $RREL_t$ is the relative bill rate; TRM_t is the term spread, the difference between the 10 year Treasury bond yield and the three month Treasury bond yield; and DEF_t is the BAA Corporate Bond rate minus the AAA Corporate Bond rate. Newey-West corrected t-statistics appear in parentheses below the coefficient estimate. Significant coefficients at the five percent level are highlighted in bold face. Regressions use data from the fourth quarter of 1952 to the fourth quarter of 2000, except for regression 9, which begins in the second quarter of 1953, the largest common sample for which all the data are available.

Table 2

Forecasting Stock Market Returns

Row	Regressors	Panel A: Forecast Horizon H in Quarters							
		1	2	4	6	8	12	16	24
1	$\widehat{cay}_t$	**1.80**	**3.06**	**5.08**	**7.02**	**7.63**	**9.94**	**10.47**	**15.44**
		(3.42)	(4.47)	(2.98)	(3.19)	(3.51)	(3.88)	(3.87)	(3.87)
		[0.07]	[0.10]	[0.14]	[0.18]	[0.17]	[0.23]	[0.22]	[0.26]
2	$d_t - p_t$	0.03	0.05	0.08	0.09	0.08	0.08	0.08	0.50
		(1.46)	(1.33)	(1.00)	(0.79)	(0.516)	(0.39)	(0.32)	(1.37)
		[0.01]	[0.01]	[0.02]	[0.02]	[0.01]	[0.00]	[0.00]	[0.10]
3	$RREL_t$	**-2.58**	**-4.14**	**-6.82**	**-6.29**	-3.19	-2.02	-2.85	-4.86
		(-3.89)	(-3.22)	(-2.79)	(-2.57)	(-1.40)	(-0.81)	(-0.86)	(-1.24)
		[0.05]	[0.06]	[0.10]	[0.06]	[0.01]	[-0.00]	[0.00]	[0.01]
4	$\widehat{cay}_t$	**1.63**	**2.80**	**4.50**	**6.51**	**7.48**	**9.91**	**10.36**	**13.08**
		(3.53)	(3.34)	(3.25)	(3.71)	(3.87)	(3.65)	(3.43)	(2.40)
	$d_t - p_t$	-0.01	-0.01	0.01	0.01	-0.00	0.02	0.00	0.40
		(-0.39)	(-0.18)	(0.10)	(0.06)	(-0.03)	(0.10)	(0.01)	(1.26)
	$RREL_t$	**-1.92**	**-3.51**	**-6.41**	**-5.26**	-1.81	0.13	-0.92	3.99
		(-2.20)	(-2.50)	(-2.98)	(-2.57)	(-0.85)	(0.04)	(-0.28)	(1.22)
		[0.10]	[0.13]	[0.20]	[0.20]	[0.16]	[0.22]	[0.21]	[0.32]
Row	Regressors	Panel B: Forecast Horizon H in Months							
		1	3	6	12	24	48	72	84
5	$\widehat{cay}_t$	**0.52**	**1.54**	**2.76**	**4.96**	**8.77**	**12.43**	**17.11**	**17.16**
		(4.15)	(4.65)	(4.08)	(4.05)	(5.89)	(9.31)	(6.81)	(5.22)
		[0.03]	[0.08]	[0.13]	[0.22]	[0.42]	[0.53]	[0.54]	[0.46]
6	lag	0.03	-0.07	-0.12	**-0.32**	**-0.85**	**-0.84**	**-1.12**	**-1.03**
		(0.51)	(-1.02)	(-1.37)	(-3.22)	(-6.73)	(-4.47)	(-5.49)	(-4.11)
	$\widehat{cay}_t$	**0.51**	**1.57**	**2.82**	**5.10**	**9.17**	**12.83**	**17.60**	**17.68**
		(4.14)	(4.67)	(4.07)	(4.03)	(6.09)	(9.80)	(7.19)	(5.47)
		[0.03]	[0.08]	[0.13]	[0.23]	[0.46]	[0.55]	[0.56]	[0.47]

Notes: See next page.

Notes to Table 2

The table reports results from long-horizon regressions of excess returns on lagged variables. H denotes the return horizon in quarters. Data in Panel A are quarterly; in Panel B data are monthly. The regressors are as follows: lag, which denotes a one-period lag of the dependent variables, one-period lagged values of the deviations from trend $\widehat{cay}_t = c_t - \widehat{\beta}_a a_t - \widehat{\beta}_y y_t$, the log dividend yield $d_t - p_t$, the dividend earnings ratio $d_t - e_t$, the detrended short-term interest rate $RREL_t$, and combinations thereof. For each regression, the table reports OLS estimates of the regressors, Newey-West corrected t-statistics in parentheses and adjusted R^2 statistics in square brackets. Significant coefficients at the five percent level are highlighted in bold. The sample period is fourth quarter of 1952 to fourth quarter of 2000 for quarterly forecasts; for monthly forecasts, the sample period is January 1959 to December 2000.

Table 3

VAR Stock Market Returns Regressions

Row	Implied R^2 for Forecast Horizon H in Quarters							
	1	2	4	6	8	12	16	24
1	0.08	0.14	0.21	0.24	0.25	0.24	0.22	0.16
	Implied R^2 for Forecast Horizon H in Months							
	1	3	6	12	24	48	72	84
2	0.04	0.09	0.16	0.24	0.31	0.28	0.23	0.20

The table reports implied R^2 statistics for H-period stock market returns from bivariate vector autoregressions (VARs) for $r_t - r_{f,t}$, the log of excess returns for the CRSP Value-Weighted stock market index, and $\widehat{cay}_t \equiv c_t - \widehat{\alpha}_a a_t - \widehat{\alpha}_y y_t$. One lag is used in the quarterly VARs; three lags are used in the monthly VARS. The implied R^2 statistics for stock market returns for horizon H are calculated from the estimated parameters of the VAR and the estimated covariance matrix of VAR residuals. The sample period is fourth quarter of 1952 to fourth quarter of 2000 for quarterly forecasts; for monthly forecasts, the sample period is January 1959 to December 2000.

Table 4

Forecasting Stock Market Volatility

Row	Regressors	Panel A: Forecast Horizon H in Quarters							
		1	2	4	6	8	12	16	24
1	V_t	**0.36**	**0.41**	**0.37**	0.21	0.049	-0.07	-0.76	-0.03
		(5.07)	(4.32)	(2.25)	(1.59)	(0.48)	(-0.71)	(-0.62)	(-0.20)
	V_{t-1}	**0.21**	0.10	-0.08	-0.13	-0.07	-0.02	0.04	0.26
		(3.04)	(1.12)	(-0.78)	(-0.96)	(0.63)	(-0.13)	(0.28)	(1.40)
		[0.22]	[0.19]	[0.11]	[0.04]	[-0.00]	[-0.01]	[-0.00]	[0.03]
2	$\widehat{cay}_t$	**-0.78**	**-1.13**	**-1.59**	**-1.84**	**-1.79**	**-1.12**	-0.31	-0.48
		(-3.90)	(-4.27)	(-4.18)	(-4.03)	(-3.54)	(-2.05)	(-0.50)	(-0.79)
		[0.12]	[0.15]	[0.18]	[0.19]	[0.15]	[0.04]	[-0.00]	[0.00]
3	V_t	**0.28**	**0.35**	**0.30**	0.14	0.02	-0.06	-0.06	0.05
		(3.76)	(4.01)	(2.07)	(1.41)	(0.18)	(-0.55)	(-0.46)	(0.27)
	V_{t-1}	**0.21**	0.07	-0.09	-0.07	0.06	0.09	0.08	0.30
		(3.59)	(1.00)	(-0.89)	(-0.54)	(0.49)	(0.43)	(0.48)	(1.54)
	$\widehat{cay}_t$	**-0.54**	**-0.88**	**-1.42**	**-1.74**	**-1.86**	**-1.26**	-0.41	-1.04
		(-3.55)	(-4.16)	(-4.32)	(-3.98)	(-3.56)	(-1.99)	(-0.57)	(-1.53)
		[0.26]	[0.28]	[0.24]	[0.20]	[0.14]	[0.04]	[-0.00]	[0.05]
4	V_t	**0.33**	**0.39**	**0.36**	0.22	0.09	-0.02	-0.07	-0.09
		(4.54)	(4.35)	(2.35)	(1.89)	(0.91)	(-0.25)	(-0.60)	(-0.54)
	V_{t-1}	**0.21**	0.10	-0.06	-0.08	-0.01	0.00	0.05	0.25
		(3.47)	(1.22)	(-0.61)	(-0.62)	(-0.06)	(0.02)	(0.29)	(1.30)
	$d_t - p_t$	**-0.01**	**-0.02**	**-0.04**	**-0.04**	-0.05	-0.03	-0.01	0.05
		(-2.35)	(-2.69)	(-2.43)	(-2.05)	(-1.56)	(-0.86)	(-0.16)	(1.25)
		[0.24]	[0.23]	[0.16]	[0.09]	[0.04]	[0.01]	[-0.01]	[0.05]
5	V_t	**0.27**	**0.34**	**0.30**	0.15	0.05	-0.02	-0.06	-0.00
		(3.64)	(3.99)	(2.11)	(1.54)	(0.49)	(-0.19)	(-0.48)	(-0.02)
	V_{t-1}	**0.21**	0.07	-0.08	-0.05	0.09	0.11	0.08	0.30
		(3.72)	(1.06)	(-0.79)	(-0.41)	(0.80)	(0.47)	(0.47)	(1.45)
	$\widehat{cay}_t$	**-0.47**	**-0.75**	**-1.26**	**-1.57**	**-1.71**	-1.21	-0.40	-1.27
		(-3.35)	(-3.73)	(-3.82)	(-3.49)	(-3.41)	(-1.97)	(-0.56)	(-1.54)
	$d_t - p_t$	-0.01	-0.01	-0.02	-0.02	-0.03	-0.03	-0.00	0.06
		(-1.29)	(-1.49)	(-1.26)	(-1.03)	(-0.97)	(-0.65)	(-0.09)	(1.39)
		[0.26]	[0.28]	[0.25]	[0.21]	[0.15]	[0.05]	[-0.01]	[0.09]

Table 4 continued, next page.

Table 4 (continued)

Forecasting Stock Market Volatility

Row	Regressors	Panel A (Continued): Forecast Horizon H in Quarters							
		1	2	4	6	8	12	16	24
6‡	V_t	0.14	**0.20**	0.13	-0.05	-0.18	-0.23	**-0.29**	-0.21
		(1.84)	(2.05)	(0.95)	(-0.46)	(-1.24)	(-1.79)	(-2.51)	(-1.48)
	V_{t-1}	**0.16**	0.02	-0.14	-0.08	0.09	0.08	-0.02	0.10
		(2.91)	(0.22)	(-1.25)	(-0.55)	(0.90)	(0.43)	(-0.11)	(0.78)
	$\widehat{cay}_t$	**-0.40**	**-0.65**	**-1.14**	**-1.42**	**-1.50**	-0.90	0.20	0.23
		(-2.85)	(-3.39)	(-3.72)	(-3.31)	(-3.28)	(-1.63)	(0.31)	(0.38)
	$d_t - p_t$	**-0.01**	**-0.02**	**-0.04**	**-0.04**	-0.05	-0.06	-0.04	**-0.10**
		(-2.32)	(-2.27)	(-2.32)	(-2.07)	(-1.83)	(-1.46)	(-1.17)	(-1.99)
	DEF_t	0.00	0.01	0.02	0.02	0.03	0.02	0.02	**0.06**
		(0.80)	(0.74)	(0.95)	(1.00)	(1.13)	(1.03)	(1.00)	(2.64)
	CP_t	**1.62**	**2.07**	**2.10**	**2.17**	**2.89**	1.51	0.95	0.68
		(3.39)	(3.01)	(3.02)	(2.40)	(2.83)	(1.42)	(0.82)	(0.54)
	$TB1Y_t$	0.11	0.16	0.24	0.30	0.32	**0.59**	**0.74**	**0.84**
		(1.66)	(1.48)	(1.34)	(1.43)	(1.34)	(2.35)	(2.25)	(2.37)
		[0.32]	[0.35]	[0.34]	[0.32]	[0.30]	[0.24]	[0.23]	[0.38]

Table 4 continued, next page.

Table 4 (continued)

Forecasting Stock Market Volatility

Row	Regressors	Panel B: Forecast Horizon H in Months							
		1	3	6	12	24	48	72	84
7	V_t	**0.01**	**0.41**	**0.41**	**0.40**	0.03	-0.09	-0.11	0.08
		(6.06)	(6.31)	(4.21)	(2.93)	(0.43)	(-0.97)	(-1.05)	(0.72)
	V_{t-1}	**0.20**	**0.16**	0.11	-0.13	-0.06	0.02	**0.39**	0.10
		(4.68)	(2.13)	(1.16)	(-1.38)	(-0.66)	(0.20)	(3.37)	(0.70)
		[0.31]	[0.25]	[0.20]	[0.13]	[-0.00]	[0.00]	[0.08]	[0.01]
8	V_t	**0.42**	**0.40**	**0.40**	**0.39**	0.05	-0.06	0.01	**0.24**
		(5.97)	(6.17)	(4.41)	(3.06)	(0.64)	(-0.57)	(0.10)	(1.97)
	V_{t-1}	**0.20**	**0.16**	0.10	-0.12	0.02	0.11	**0.58**	**0.31**
		(4.72)	(2.20)	(1.16)	(-1.34)	(0.19)	(0.81)	(3.97)	(2.18)
	$\widehat{cay}_t$	**-0.17**	**-0.36**	**-0.58**	**-0.98**	**-1.27**	-0.69	**-1.48**	**-1.85**
		(-2.96)	(-3.51)	(-3.77)	(-4.36)	(-3.77)	(-1.48)	(-3.90)	(-5.28)
		[0.33]	[0.28]	[0.26]	[0.22]	[0.10]	[0.02]	[0.17]	[0.15]

Notes: See next page.

Notes to Table 4

The table presents results from long-horizon regressions of stock market volatility on lagged variables using quarterly data from 1952:4-2000:4, OLS estimation. The dependent variable at each forecast horizon H is the H-step ahead volatility, equal to

$$v_{t+1,t+H} = [\Sigma_{s \in t+1,\ldots,t+h}(r_s - \bar{r})^2]^{1/2}$$

where v denotes the variance of the CRSP value-weighted index estimated from daily returns. The H-period volatilities are regressed on one-period lagged values of the log dividend yield, $d_t - p_t$, the consumption-wealth ration proxy $\widehat{cay}_t = c_t - \widehat{\beta}_a a_t - \widehat{\beta}_y y_t$, the BAA Corporate Bond rate minus the AAA Corporate Bond rate, DEF_t, the difference between the yield on six-month commercial paper and the 3-month treasury bill yield, CP_t, the one-year Treasury yield, $TB1Y_t$, and their own first and second lagged values, denoted V_t and V_{t-1}. For each regression, the table reports OLS estimates of the regressors, Newey-West corrected t-statistics in parentheses, and adjusted R^2 statistics in square brackets. Significant coefficients at the 5% level are highlighted in bold. The sample period is fourth quarter of 1952 to fourth quarter of 2000 for quarterly forecasts, except for regression 6 (marked with a ‡), which uses a sample running from the second quarter of 1953 to the fourth quarter of 2000, the largest common sample for which all the data are available. For monthly forecasts, the sample period is January 1959 to December 2000.

Table 5

Maximum Likelihood Estimates of GARCH(1,1) Model for Consumption Growth

	1	2	3	4
Mean Equation				
Constant	**0.0024**	**0.0024**	**0.0024**	**0.0027**
$(S.E.)$	(0.0006)	(0.0006)	(0.0005)	(0.0006)
Δc_{t-1}	**0.3141**	**0.2878**	**0.2721**	**0.2761**
$(S.E.)$	(0.0754)	(0.0718)	(0.0663)	(0.0685)
Δc_{t-2}	-0.0002	0.0341	0.0796	0.0420
$(S.E.)$	(0.0703)	(0.0714)	(0.0666)	(0.0753)
Δc_{t-3}	**0.2375**	**0.2168**	**0.2331**	**0.1885**
$(S.E.)$	(0.0548)	(0.0611)	(0.0537)	(0.0683)
Variance Equation				
Constant	1.38E-8	**4.68E-5**	4.05E-7	**5.28E-5**
$(S.E.)$	(1.94E-7)	(3.78E-7)	(2.07E-7)	(2.08E-7)
ϵ_{t-1}^2	-0.0175	-0.0066	-0.0206	0.0618
$(S.E.)$	(0.0290)	(0.0495)	(0.0306)	(0.0369)
σ_{t-1}^2	**1.007**	**0.0594**	**1.0037**	**0.7938**
$(S.E.)$	(0.0243)	(0.0480)	(0.0340)	(0.0517)
$\widehat{cay}_{t-1}$		**-7.64E-5**		**-8.19E-5**
$(S.E.)$		(6.04E-7)		(2.13E-7)
$r_{t-1} - r_{f,t-1}$			**-1.45E-5**	**-4.36E-5**
$(S.E.)$			(2.00E-7)	(9.95E-6)

The table reports estimates from the *GARCH(1,1)* model:

$$\Delta c_t = \alpha_0 + \alpha_1 \Delta c_{t-1} + \alpha_2 \Delta c_{t-2} + \alpha_3 \Delta c_{t-3} + \epsilon_t$$

$$\sigma_t^2 = \delta_0 + \delta_1 \epsilon_{t-1}^2 + \delta_2 \sigma_{t-1}^2 + \delta_3 X_{t-1},$$

where σ_t^2 is the conditional variance of ϵ_t. The regressors in X_{t-1} are as follows: Δc_t is consumption growth, $\widehat{cay}_{t-1} \equiv c_{t-1} - \widehat{\beta}_a a_{t-1} - \widehat{\beta}_y y_{t-1}$, and $r_{t-1} - r_{f,t-1}$ is lagged excess returns for the CRSP-VW index. Bollerslev-Wooldridge robust standard errors appear in parentheses beneath the coefficient estimates. Coefficents two standard errors or more from zero are highlighted in bold face. The sample runs from the first quarter of 1953 to the first quarter of 2001.

Table 6

Summary Statistics for Sharpe Ratios

	SR_t^{VW}	SR_t^{CC}	SR_t^{CV}
Correlation Matrix			
SR_t^{VW}	1.000	0.382	-0.312
SR_t^{CC}		1.000	-0.249
SR_t^{CV}			1.000
Univariate Summary Statistics			
Mean	0.390	0.229	0.390
Standard Deviation	0.448	0.086	0.091
Autocorrelation	0.831	0.968	0.849

SR_t^{VW} is the Sharpe Ratio estimated from the CRSP-VW index. SR_t^{CC} is the Sharpe Ratio implied by Campbell and Cochrane (1999); SR_t^{CV} is the Sharpe Ratio implied by the Consumption Volatility Model:

$$\mathrm{E}(R_{t+1}) - R_t^f = \gamma \sigma_t(\Delta c_{t+1})\ \sigma_t(R_{t+1}) \rho_t(m_{t+1}, R_{t+1})$$

where γ is the constant coefficient of risk aversion and is set equal to 92. The statistics are computed for the largest common set of available data for all the variables, which spans the fourth quarter of 1953 to the fourth quarter of 2000.

Figure 1: Conditional Volatility for the CRSP-VW Index

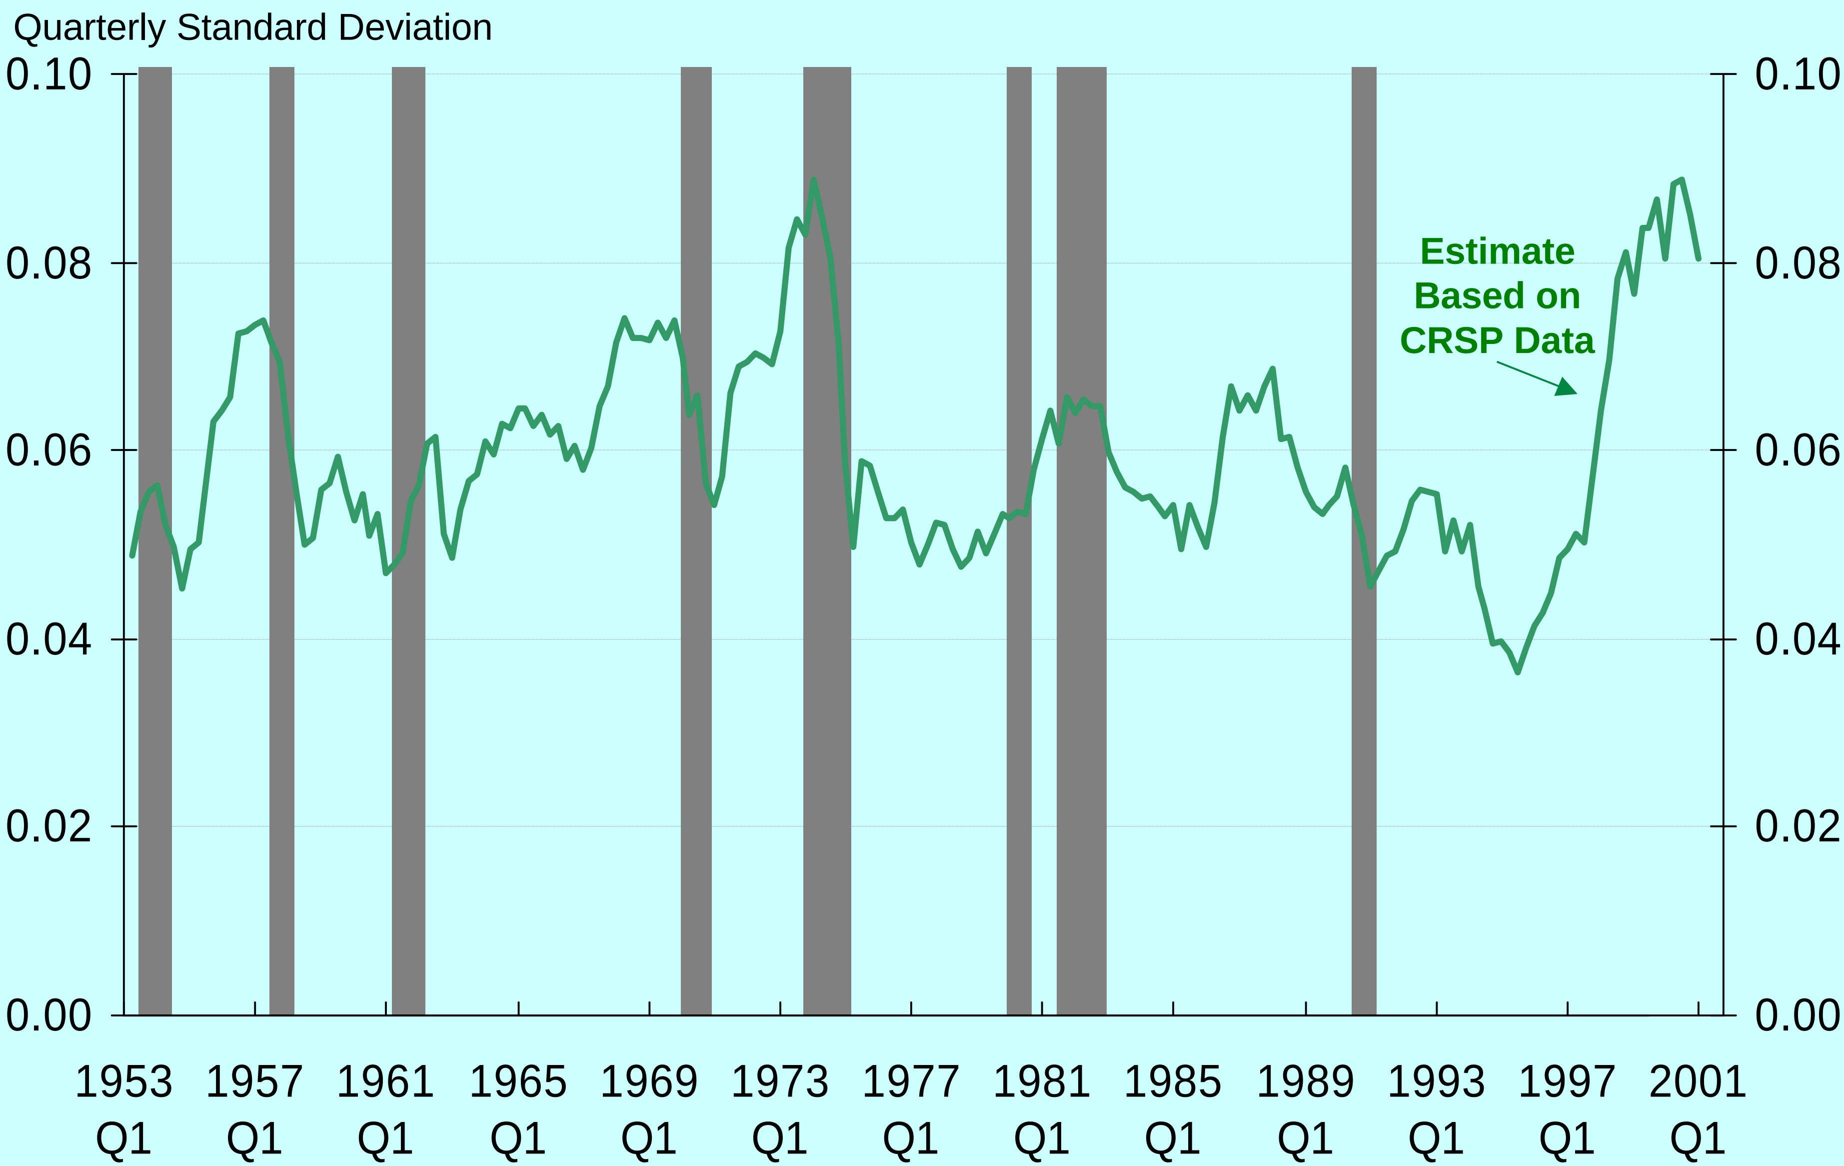

Note: Shading denotes quarters designated recession by the NBER
Source: Authors' Calculations

Figure 2: Conditional Sharpe Ratio

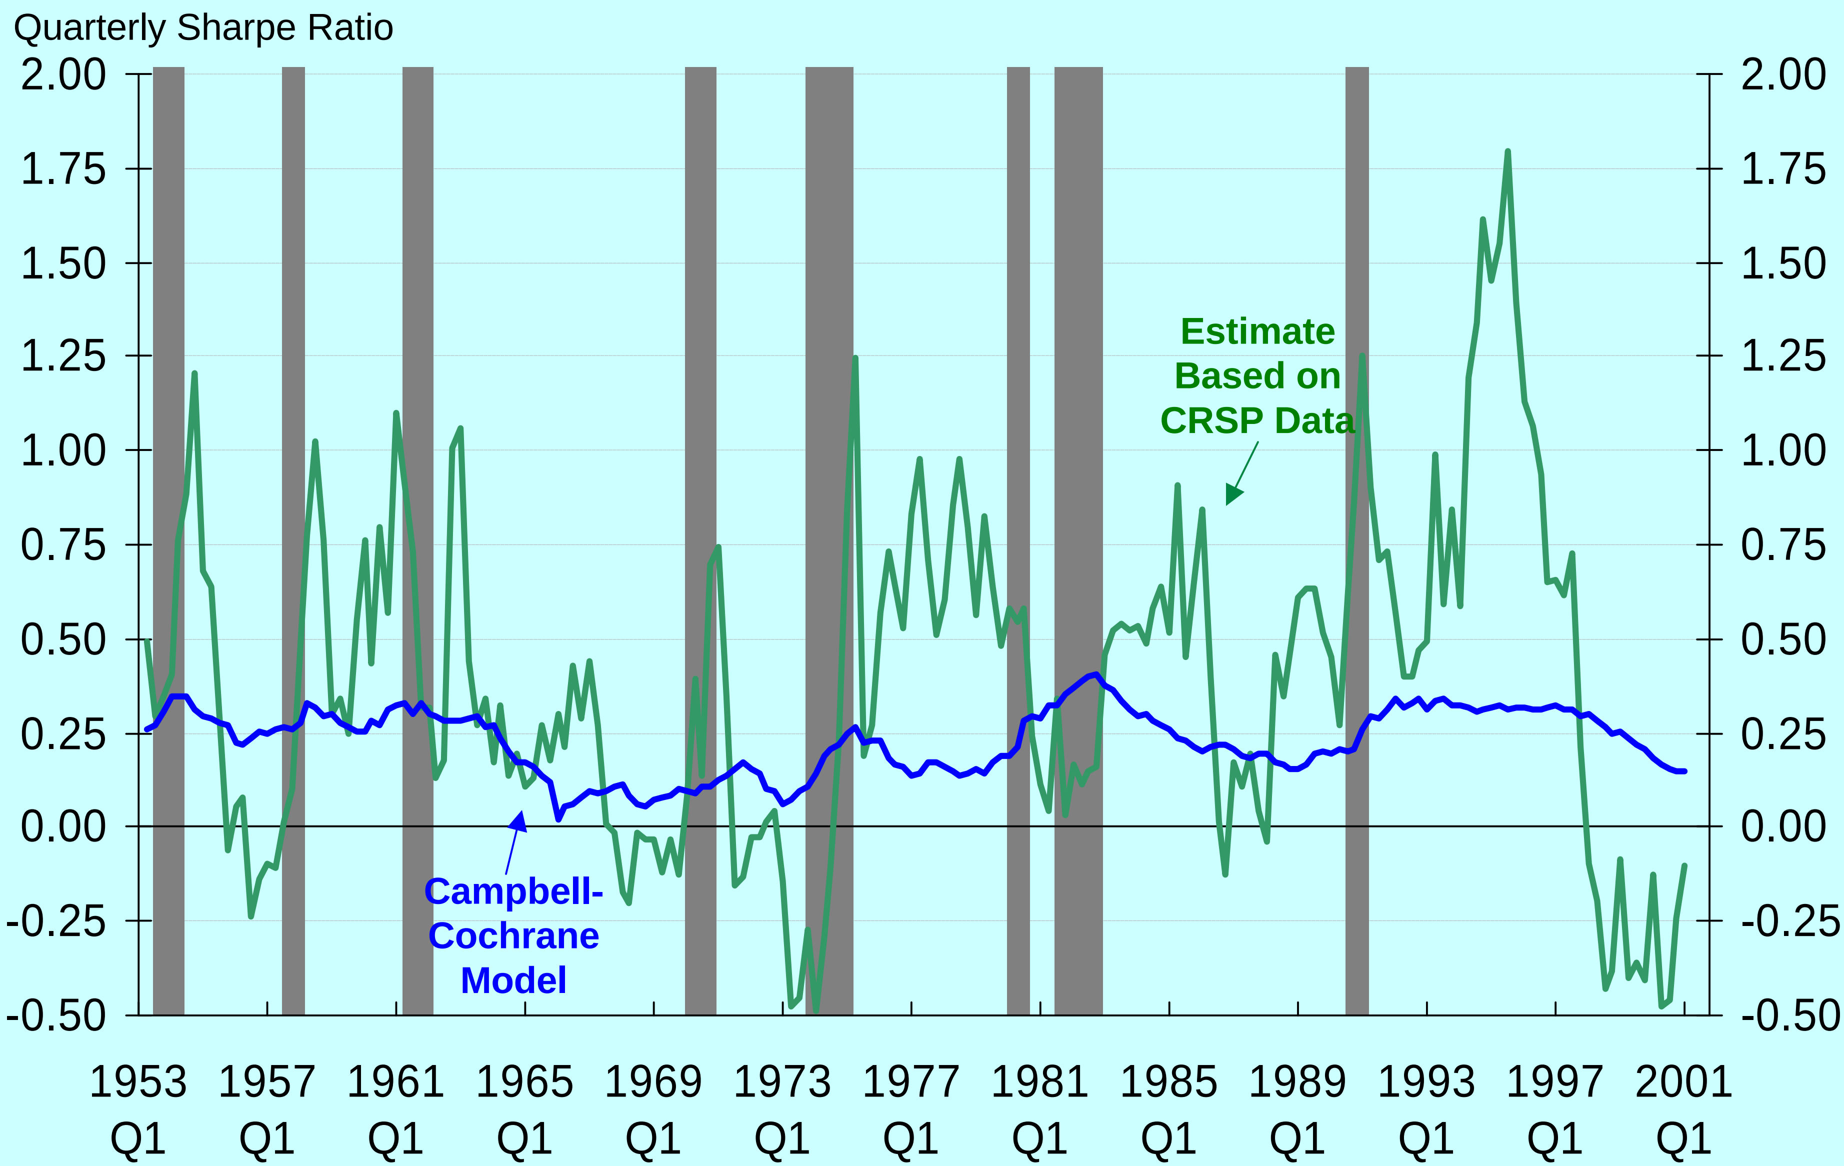

Note: Shading denotes quarters designated recession by the NBER

Sources: Authors' Calculations, Campbell and Cochrane (1999)

Figure 3: Estimates of the Sharpe Ratio from the Consumption Volatility Model and from the CRSP-VW Index

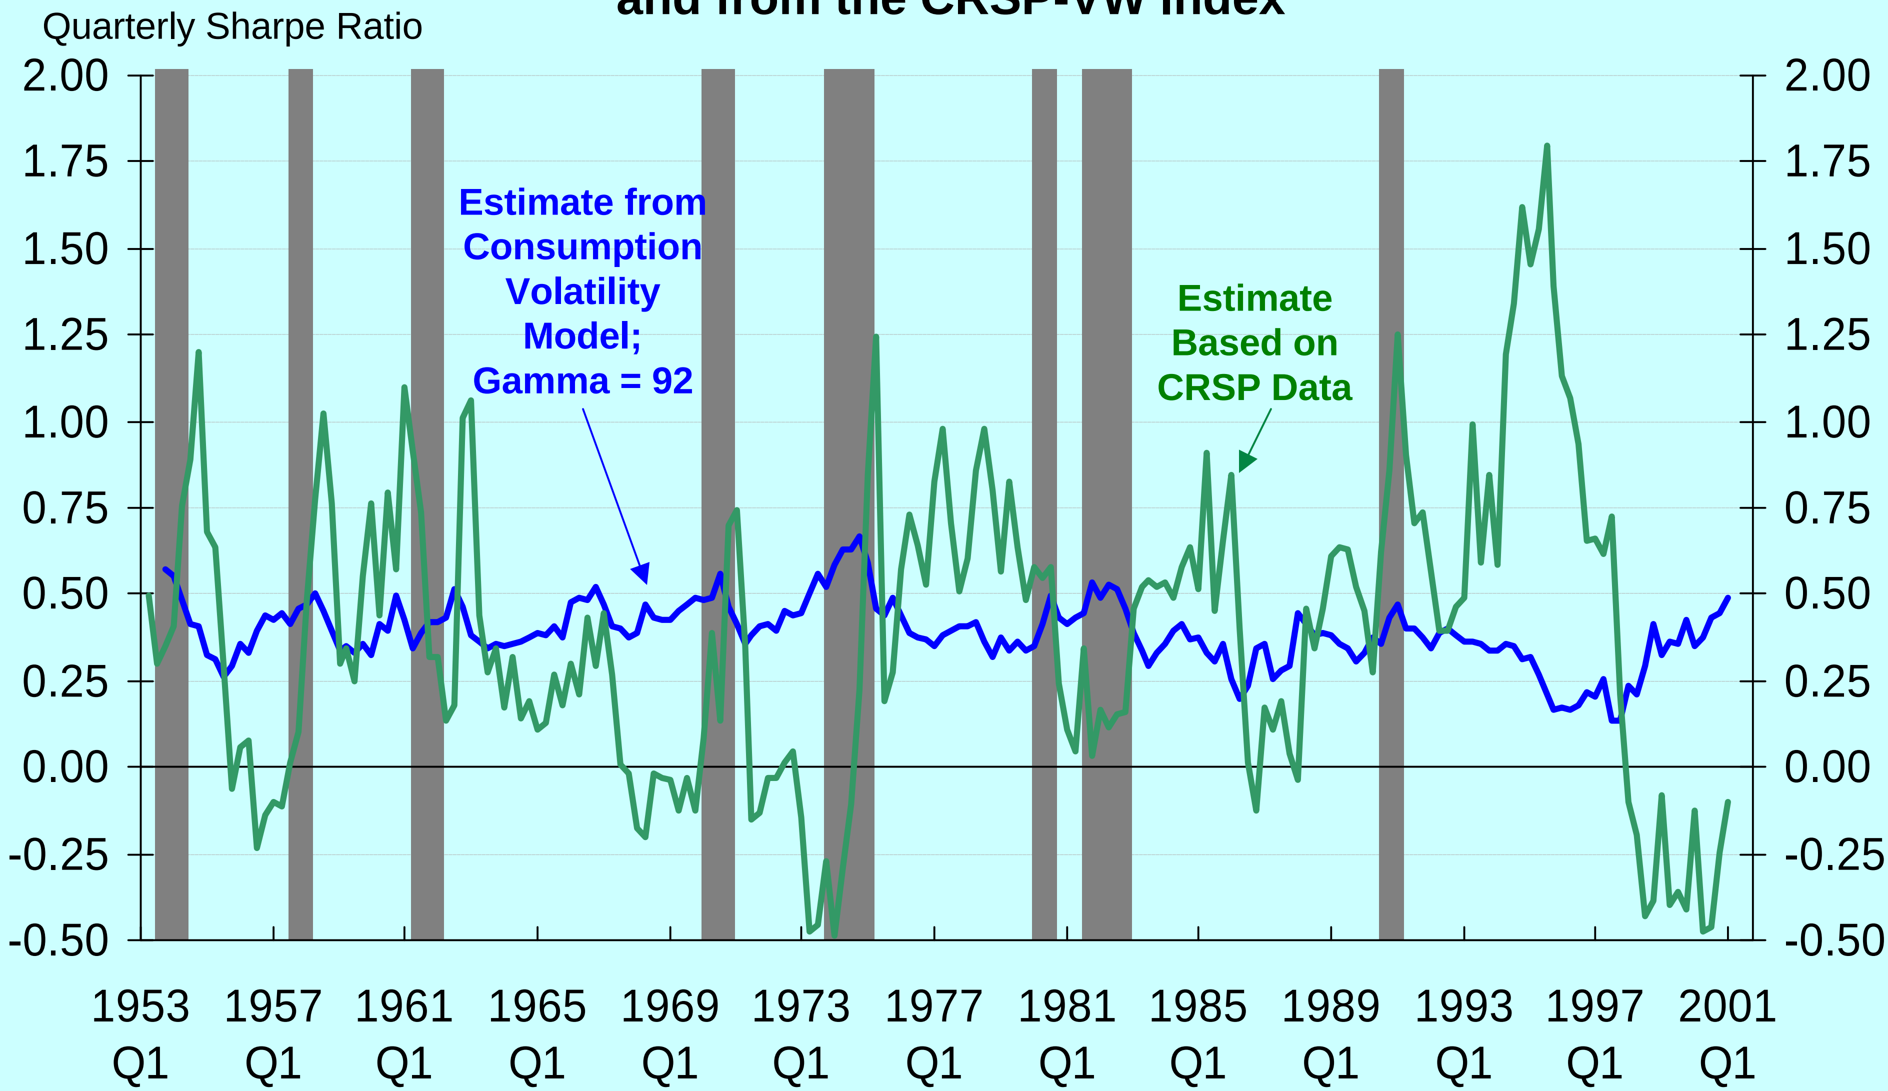

Note: Shading denotes quarters designated recession by the NBER. Gamma refers to the risk aversion scale factor in the Consumption Volatility Model. Gamma = 92 is the scale factor which equates the means of the estimates from the CRSP-VW and the Consumption Volatility Model.

Source: Authors' Calculations