Measuring the Persistence of Expected Returns

NBER WORKING PAPER SERIES

MEASURING THE PERSISTENCE OF EXPECTED RETURNS

John Y. Campbell

Working Paper No. 3305

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 1990

NBER Working Paper #3305
March 1990

MEASURING THE PERSISTENCE OF EXPECTED RETURNS

ABSTRACT

This paper summarizes earlier research on the sources of variation in monthly U.S. stock returns in the period 1927-88. A log-linear model is used to break unexpected returns into changing expectations about future dividends and changing expectations about future returns. Even though stock returns are not highly forecastable, the model attributes one-third of the variation in returns to changing expected returns, one-third to changing future dividends, and one-third to the covariance between these components. Changing expected returns have a large effect on the stock market because their movements are persistent and negatively correlated with changing expected dividends.

John Y. Campbell
London School of Economics
Financial Markets Group
Houghton Street
London WC2A 2AE
England
(011-44-1)-405-7686, x3106

A great deal of recent research has documented the fact that rational expectations of real returns on long-term financial assets move systematically through time. Most of this work concentrates on the variability of expected returns, but the persistence or serial correlation of expected returns is also an important issue. If expected returns follow a persistent time-series process, then movements in expected returns will have a large impact on asset prices; prices are much less sensitive to transitory fluctuations in expected returns. Thus any attempt to explain the variability of asset prices or returns requires information on the persistence of movements in expected returns.

Information on persistence can also guide the search for an economic explanation of movements in expected returns. Variable expected returns need to be accounted for by variable intertemporal marginal rates of substitution, either those of a representative agent in a standard finance-theoretic model, or those of optimizing agents in a model with "noise traders". The time-series properties of expected returns can be used to restrict the time-series properties of the relevant intertemporal marginal rate of substitution.

In this paper I estimate the persistence of expected returns in the U.S. stock market, and use the estimates to account for the variability of unexpected stock returns. I begin with a standard regression model which forecasts monthly stock returns with a modest R^2 statistic of less than 7%. I incorporate this model into a vector autoregression which

describes the evolution through time of the forecasting variables, and thus the expected stock return. I then calculate the effect on the stock price of an innovation in the expected return; this effect is large because movements in the expected return are persistent. Finally, I use the system to decompose the variance of the unexpected stock return into three components: the variance of changes in rational expectations of future returns, the variance of changes in rational expectations of future dividends, and a covariance term. The variance of changes in expectations of future returns and the covariance term are always important components of the overall variance of stock returns.

Another way to express this result is that the overall variance of stock returns is always greater than the variance of news about cash flows. Short-term predictability of returns can increase the variance of unexpected returns to a surprising degree. The findings here suggest that a satisfactory explanation of stock market volatility must simultaneously explain the short-term predictability of stock returns.

I. Expected returns and unexpected returns

In order to characterize the relation between changing expected returns and unexpected returns, I will use the log-linear "dividend-ratio model" developed in my work with Robert Shiller (1988a, 1988b). This model is an appropriate framework because it relates asset prices to both expected returns and expected future cash flows. The model says that, to a first-order Taylor approximation,

$$(1) \quad h_{t+1} \approx k + \delta_t - \rho\delta_{t+1} + \Delta d_{t+1},$$

where h_{t+1} is the log holding-period return on a stock held from the end of period t to the end of period t+1, δ_t is the log dividend-price ratio $d_t - p_t$, d_t is the log dividend paid during period t, p_t is the log stock price at the end of period t, and ρ and k are parameters of linearization. The parameter ρ is the average ratio of the stock price to the sum of the stock price and the dividend, a number a little smaller than one.[1] Equation (1) says that the return on stock is high if the dividend-price ratio is high when the stock is purchased, if dividend growth occurs during the holding period, and if the dividend-price ratio falls during the holding period.

Equation (1) can be thought of as a difference equation relating δ_t to δ_{t+1}, Δd_{t+1} and h_{t+1}. Solving forward, and imposing the terminal condition that $\lim_{i\to\infty} \rho^i \delta_{t+i} = 0$, one obtains

$$(2) \qquad \delta_t = \sum_{j=1}^{\infty} \rho^{j-1}[h_{t+j} - \Delta d_{t+j}] - k/(1-\rho).$$

This equation says that a high dividend-price ratio today must generate high future returns unless dividend growth is low in the future. It is important to note that all the variables in (2) are measured ex post; (2) has been obtained <u>only</u> by the linear approximation of h_t and the imposition of a condition that δ_{t+i} does not explode as i increases.

However (2) also holds ex ante. If one takes rational expectations of equation (2), conditional on information available at the end of time period t, the left hand side is unchanged since δ_t is in the information set, and the right hand side becomes an <u>expected</u> discounted value. Using the ex ante version of (2) to substitute δ_t and δ_{t+1} out of (1), I obtain

$$(3)\quad h_{t+1} - E_t h_{t+1} = (E_{t+1}-E_t)\sum_{j=1}^{\infty}\rho^{j-1}\Delta d_{t+j} - (E_{t+1}-E_t)\sum_{j=1}^{\infty}\rho^{j}h_{t+1+j},$$

or in more compact notation,

$$v_{t+1} = \eta_{d,t+1} - \eta_{h,t+1}.$$

Equation (3) makes the central point of this paper. If expectations are internally consistent, then an unanticipated capital gain on the stock must be associated with either an increase in expected future dividends $\eta_{d,t+1}$ or a fall in expected future returns $-\eta_{h,t+1}$. To see why this is so, consider an asset with fixed dividends whose price rises. This asset's dividend yield is now lower, so its returns must be lower at some point in the future unless the asset price follows an explosive upward path which I have ruled out by assumption.

Equation (3) can be used to decompose unexpected stock returns into the two components $\eta_{d,t+1}$ and $-\eta_{h,t+1}$. In what follows, I shall refer to the former as "news about future dividends", and to the latter as "news about future expected returns". This should not be taken to imply one-way causality from expected dividends and returns to prices; in general these variables are all determined simultaneously.[2]

The importance of persistence can be most easily appreciated by working out in detail the special case in which the expected return follows a first-order univariate autoregressive, or AR(1) process with coefficient ϕ. I define u_{t+1} to be the innovation at time t+1 in the one-period-ahead expected return, $u_{t+1} = (E_{t+1}-E_t)h_{t+2}$. In the AR(1) case, u_{t+1} satisfies

$$(4)\quad E_{t+1}h_{t+2} = \phi E_t h_{t+1} + u_{t+1}.$$

It is important to note that equation (4) does not restrict the size of the market's information set. In particular, there is no presumption that the relevant information set contains only the history of past asset returns. It is quite possible that a very large number of variables is useful in forecasting the asset return over the next period; equation (4) merely restricts the way in which the next period's forecast is related to past forecasts.

Equation (4) implies that $\eta_{h,t+1} = \rho u_{t+1}/(1-\rho\phi)$. Since ρ is a number very close to one, this means that a 1% increase in the expected return today is associated with a capital loss of about 2% if the AR coefficient is 0.5, a loss of about 4% if the AR coefficient is 0.75, and a loss of about 10% if the AR coefficient is 0.9. James Poterba and Lawrence Summers (1988) give a similar result.

One can also calculate the ratio of the variance of news about future expected returns to the total variance of unexpected returns. Equations (3) and (4) imply that for ρ close to one, $\mathrm{Var}(\eta_{h,t+1})/\mathrm{Var}(v_{t+1}) \approx ((1+\phi)/(1-\phi))(R^2/(1-R^2))$, where R^2 is the fraction of the variance of stock returns which is predictable. If R^2 is 0.025, then the share of news about future expected returns in the variance of unexpected returns is 0.08 for $\phi = 0.5$, 0.18 for $\phi = 0.75$, and a startling 0.49 for $\phi = 0.9$. These parameters are not unreasonable ones for monthly stock returns. An R^2 of 0.025 is quite modest, and a process with $\phi = 0.9$ has a half-life of only a little more than six months. The lesson is that apparently trivial but persistent movements in expected returns can be a major force driving unexpected returns. The R^2 for one-period returns is thus an inadequate measure of the importance of expected return variation.

The assumption that the expected stock return follows a univariate AR(1) is highly restrictive. Fortunately it is possible to generalize the previous discussion to handle the case where the expected return is one element of a vector autoregression (VAR). Shmuel Kandel and Robert Stambaugh (1989) also propose a VAR to describe stock returns.

First, I define a vector z_{t+1} which has k elements, the first of which is the stock return h_{t+1}. The other elements are other variables which are known to the market by the end of period t+1. Then I assume that the vector z_{t+1} follows a first-order VAR:

$$(5) \quad z_{t+1} = Az_t + w_{t+1}.$$

The assumption that the VAR is first-order is not restrictive, since a higher-order VAR can always be stacked into first-order (companion) form in the manner discussed in my paper with Shiller (1988a). The matrix A is known as the companion matrix of the VAR.

The first-order VAR generates simple multi-period forecasts of returns. Define el to be a k-element vector, whose first element is 1 and whose other elements are all 0. This vector picks out the stock return from the variables of the VAR: thus the unexpected stock return $v_{t+1} = el'w_{t+1}$. Expected future stock returns are given by $E_t h_{t+1+j} = el'A^{j+1}z_t$. It follows that the discounted sum of revisions in forecast returns can be written as

$$(6) \quad \eta_{h,t+1} = el' \sum_{j=1}^{\infty} \rho^j A^j w_{t+1} = \lambda' w_{t+1},$$

where λ' is defined to equal $\rho el'(I-\rho A)^{-1}A$, a nonlinear function of the VAR coefficients. From equation (3), $\eta_{d,t+1}$ equals $(el'+\lambda')w_{t+1}$. These expressions can be used to decompose the variance of the unexpected stock return, v_{t+1}, into the variance of the component which is due to expected return variation, $\eta_{h,t+1}$, the variance of the news about dividends, $\eta_{d,t+1}$, and a covariance term.

In the VAR context there is no single measure of the persistence of expected returns. But one natural way to summarize persistence is by the variability of the innovation in the expected present value of future returns, relative to the variability of the innovation in the one-period-ahead expected return. Thus I define the VAR persistence measure P as

$$(7) \quad P \equiv \sigma(\eta_{h,t+1})/\sigma(u_{t+1}) = \sigma(\lambda' w_{t+1})/\sigma(el'Aw_{t+1}),$$

where $\sigma(x)$ denotes the standard deviation of x. Another way to describe the statistic P is to say that a typical 1% innovation in the expected return will be associated with a P% capital loss on the stock. In the univariate AR(1) case discussed earlier, $P = \rho/(1-\rho\phi) \approx 1/(1-\phi)$.

II Application to the U.S. stock market

For the sake of comparability with previous work, I use a standard data set here. I study the behavior of the monthly value-weighted New York Stock Exchange Index, as reported by the Center for Research in Security Prices (CRSP) at the University of Chicago. The data set runs from 1926 to 1988, but I reserve the first year for lags so that my full sample period is 1927 to 1988. I deflate the nominal return on the index

using the Consumer Price Index reported in Ibbotson Associates (1989).

The forecasting variables I use for the stock return are the lagged stock return, the dividend-price ratio D/P, and the "relative bill rate" RREL, the difference between a short-term Treasury bill rate and its one-year backward moving average. The lagged stock return is included because the stock return forecasting equation will be one equation of a VAR system. The dividend-price ratio is included because the ex ante version of equation (2) shows that it will reflect any changes that may occur in future expected returns. (See also Eugene Fama and Kenneth French, 1988.) The ratio is measured as total dividends paid over the previous year, divided by the current stock price.

The relative bill rate is included because many authors have noted that the level of short-term interest rates helps to forecast stock returns. The short-term interest rate itself may be nonstationary over this sample period, so it needs to be stochastically detrended. The subtraction of a one-year moving average is a crude way to do this; the relative bill rate can also be written as a triangular moving average of changes in the short-term interest rate, so it is stationary in levels if the short rate is stationary in differences.[3] The short rate used is the one-month Treasury bill rate series from Ibbotson Associates (1989).

One problem which arises when interest rate data are used is that the behavior of interest rates has changed over time. In particular, the Federal Reserve Board held interest rates almost constant for much of the period up to 1951, when a Federal Reserve Board-Treasury Accord allowed rates to move more freely. Accordingly, I split the 1926-88 sample at the end of 1951. This also allows a separate look at the data from the

period around the Great Depression, which may behave differently from the postwar data (Myung Jig Kim, Charles Nelson, and Richard Startz 1989).

The stock return forecasting equation estimated over the period 1927-88 is $h_{t+1} = 0.107\ h_t + 0.331\ (D/P)_t - 0.424\ RREL_t$. The equation has an R^2 of 0.024, and the forecasting variables are jointly significant at the 1.8% level using a heteroskedasticity-consistent test statistic. When D/P is regressed on the same forecasting variables, its own lag receives a coefficient of 0.963 while the other variables have almost zero coefficients. RREL behaves in a similar fashion with an own lag coefficient of 0.669.

These estimates have quite striking implications for the variance decomposition of unexpected stock returns. The variance of news about future expected returns $\eta_{h,t+1}$ accounts for 0.28 of the total variance of unexpected stock returns, the variance of news about future dividends $\eta_{d,t+1}$ accounts for 0.37, while the covariance of these two terms accounts for the remaining 0.35. The covariance term is large because news about future returns and news about future dividends have a negative correlation of -0.53. When the stock market rises because of good news about future dividends, expected returns fall and the market rises even further.

Shocks to expected returns in this system can come from any of the three variables in the VAR. Each type of shock has a different persistence. The average persistence measure P is estimated to be 4.8, indicating that a typical 1% shock to the expected return is associated with a capital loss of almost 5%.

When the VAR system is estimated for the period 1927-51, the

forecasting variables are jointly significant for stock returns at only the 18% level. The problem is partly that stock returns have a high variance which moves with the level of the dividend-price ratio; this means that heteroskedasticity correction greatly increases the standard error on this variable. Unsurprisingly, given the interest rate regime, the interest rate variable is a poor forecaster before 1952. Nonetheless, even in this period the variance decomposition of unexpected returns attributes only 0.44 of the variance to the variance of news about dividends. The remainder is attributed to the variance of news about expected returns (0.19) or the covariance term (0.38).

In 1952-88 stock returns are more strongly forecastable; the forecasting system used here does <u>not</u> derive its predictive power from the Great Depression period. The returns equation has an R^2 of 6.5%, and both the dividend-price ratio and the relative bill rate are highly significant. 0.77 of the variance of unexpected stock returns is now attributed to news about future expected returns, and only 0.13 to news about dividends. Once again the covariance term is positive because there is a negative correlation between news about dividends and news about expected returns.

The results presented here, and in greater detail in my 1990 paper, suggest that predictable time variation in stock returns is extremely important for understanding the volatility of the stock market. Predictable returns move quite persistently, and they tend to fall when expected dividends rise, amplifying the response of the stock market to news about future dividends. These facts pose a challenge to asset pricing theory.

Bibliography

Campbell, John Y., "A Variance Decomposition for Stock Returns", NBER Working Paper No. 3246, January 1990.

Campbell, John Y. and Robert J. Shiller (1988a), "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors", Review of Financial Studies, Fall 1988, 1, 195-228.

Campbell, John Y. and Robert J. Shiller (1988b), "Stock Prices, Earnings, and Expected Dividends", Journal of Finance, July 1988, 43, 661-676.

Fama, Eugene F. and Kenneth R. French, "Dividend Yields and Expected Stock Returns", Journal of Financial Economics, October 1988, 22, 3-25.

Ibbotson Associates, Stocks, Bonds, Bills, and Inflation: 1989 Yearbook, Chicago, 1989.

Kandel, Shmuel and Robert F. Stambaugh, "Modelling Expected Stock Returns for Long and Short Horizons", unpublished paper, Wharton School, University of Pennsylvania, July 1989.

Kim, Myung Jig, Charles R. Nelson, and Richard Startz, "Mean Reversion in Stock Prices? A Reappraisal of the Empirical Evidence", unpublished paper, University of Washington, May 1989.

Poterba, James M. and Lawrence H. Summers, "Mean Reversion in Stock Prices: Evidence and Implications", Journal of Financial Economics, October 1988, 22, 27-59.

Footnotes

*. Woodrow Wilson School, Princeton University, Princeton NJ 08544. I am grateful to the National Science Foundation and the Sloan Foundation for financial support, and to Robert Shiller for exceptional assistance.

1. Equation (1) and the other formulas given here differ slightly from those given in my papers with Shiller because the notation here uses a different timing convention. In this paper I define the time t stock price and conditional expectation of future variables to be measured at the end of period t rather than the beginning of period t. This conforms with the more standard practice in the finance literature.

2. Consider for example a model in which noise traders become bullish on stock, driving prices up and rational expectations of future returns down. In this model, the news is simply the innovation in the noise traders' demand for stock. Nonetheless (3) must still describe the expectations of any rational agent.

3. Another recently popular way to detrend the interest rate is to use the yield spread between interest rates of two different maturities. The relative bill rate has at least as much forecasting power for stock returns as the long-short yield spread, which is insignificant when it is added to the equations reported below.