

may be no premium at all! It is important to note, however, that each of these papers attributes equal weight to all stocks that are used in the calculation of the mean or median estimate of the market expected rate of return in Claus and Thomas (2001) and Gebhardt, Lee, and Swaminathan (1999), and in the regression in Easton, Taylor, Shroff, and Sougiannis (2002).

This equal-weighting has two potential effects. First, small stocks have an undue effect on the estimate of the market return. Second, stocks with low or negative earnings, which are somewhat meaningless as summary valuation metrics, potentially have an influence that is similar to the influence of large stable firms where earnings are a much more meaningful valuation metric. In order to avoid these undue influences, we repeat all of the analyses weighting each of the observations by market capitalization.

Our estimate of the implied expected rate of return on the market from the value-weighted regression, after removing the effect of bias in analysts' forecasts, is 9.67 percent with an implied equity premium of 4.43 percent. Of course, this estimate of the equity premium is more reasonable than that obtained when all observations have equal weight. We also find that the extent of analysts' optimism decreases as firm size increases. The effect of analysts' bias on the estimate of the implied expected rate of return on the market that is based on the value-weighted regression is lower than the estimate from the equally-weighted regression; 1.60 percent compared with 2.84 percent.

Studies such as Michaely and Womack (1999); Boni and Womack (2002); Eames, Glover, and Kennedy (2002); and Bradshaw (2004) show that analysts generally make "strong buy" and "buy" recommendations. They sometimes recommend "hold", and rarely recommend "sell". It seems reasonable to expect that buy recommendations will be associated with ex ante

optimistic forecasts. In other words, the pervasiveness of buy recommendations may explain the optimistic bias in forecasts and in expected rates of return based on analysts' forecasts.

To examine this issue further, we repeat the analyses for sub-samples formed on the basis of number of analysts comprising the consensus who recommend "buy". Contrary to our expectations, we show that the consensus analyst forecast is optimistic even when less than 30 percent of analysts' comprising the consensus recommend "buy".⁶ Estimates of the implied expected rate of return are biased upward even for these sub-samples. Interestingly, we show that the implied expected rate of return declines monotonically as the percentage of analysts recommending "buy" declines. In other words, analysts' recommendations appear to be based on expected rates of return rather than the difference between the analysts' expectations and the market expectation. This evidence is consistent with the observation in Groysberg, Healy, Chapman, and Gui (2006) that analysts' salary increases and bonuses are based on stock returns subsequent to their recommendations adjusted for the return on the S&P 500 index.

The remainder of the paper proceeds as follows. In section 2, we outline the methods used in estimating the expected rate of return implied by market prices, current book value of equity, and current and forecasted accounting earnings. Section 3 describes the data used in our analyses. In section 4, we document the ex post and the ex ante bias in consensus analysts' forecasts and discuss the implications for cost of capital estimates in extant accounting research, which are generally based on equal weighting of observations from the entire sample of firms followed by analysts. In section 5, we repeat the analyses using value-weighting of firms to show that the estimate of the bias is lower and the estimate of the expected equity risk premium is more reasonable than that obtained in extant studies. Sub-samples based on percentage of

⁶ While it is reasonable to expect that the level of the analyst's recommendation should be associated with *expected* abnormal returns, it should be noted that Bradshaw (2004) finds analysts' recommendations uncorrelated with future *realized* abnormal returns.

analysts recommending buy are analyzed in section 6. Section 7 concludes with a summary of implications for future research.

2. Methods of estimating the implied expected rate of return

We develop three methods for estimating the implied expected rate of return. These estimates, which are based on (1) I/B/E/S earnings forecasts, (2) realized earnings, and (3) perfect foresight forecasts of earnings, lead to two determinations of the bias when estimates of the market expected rate of return are based on analysts' forecasts of earnings. Each of these methods determines bias as the difference between estimates based on forecasts of earnings and estimates based on earnings realizations.

We refer to the primary measure as the *ex ante* measure of bias because it relies on information available at the time of the earnings forecast. This measure compares the estimates of the implied expected rate of return based on analysts' forecasts with estimates based on current earnings realizations. The other measure compares estimates formed using analysts' forecasts with estimates based on perfect foresight of next-period earnings realizations. We refer to this as the *ex post* measure. We note there may be factors other than analysts' optimism affecting each of these measures of bias; but, since other factors affecting the *ex ante* measure would not affect the *ex post* measure (and vice-versa), obtaining similar results based on both measures suggests that the effect of other factors is minimal. We elaborate on this point in section 2.3.

2.1. Ex ante determination of the effect of bias

Each of the methods for estimating the implied expected rate of return are derived from the residual income valuation model which may be written as follows:

$$v_{jt} \equiv bps_{jt} + \sum_{\tau=1}^{\infty} \frac{eps_{jt+\tau} - r_j \times bps_{jt+\tau-1}}{(1+r_j)^\tau} \quad (1)$$

where v_{jt} is the intrinsic value per share of firm j at time t , bps_{jt} is the book value per share of common equity of firm j at time t , eps_{jt} is the earnings per share of firm j at time t and r_j is the cost of capital for firm j .⁷ Easton, Taylor, Shroff, and Sougiannis (2002) rely on the following finite horizon version of this model:

$$p_{jt} \equiv bps_{jt} + \frac{eps_{jt+1}^{IBES} - r_j \times bps_{jt}}{(r_j - g_j)} \quad (2)$$

where p_{jt} is price per share for firm j at time t , eps_{jt+1}^{IBES} is an I/B/E/S forecast of earnings for period $t+1$, and g_j is the expected rate of growth in residual income beyond period $t+1$ required to equate $(p_{jt} - bps_{jt})$ and the present value of an infinite residual income stream.^{8,9}

Easton, Taylor, Shroff, and Sougiannis (2002), like many other studies, implicitly use analysts' forecasts of earnings as a proxy for market expectations of next period earnings. Optimistic bias in analysts' forecasts implies a bias in this proxy. In this paper we use a modification of the method in O'Hanlon and Steele (2000) to determine, ex ante, the effect of the forecast error on the estimate of the expected rate of return. This method provides an estimate of the expected rate of return implied by current realized accounting earnings; we compare this with

⁷ Derivation of this model requires the no arbitrage assumption, which is necessary to derive the dividend capitalization formula, and that earnings are comprehensive – in other words, the articulation of earnings and book value is clean surplus.

⁸ Price in this relation replaces intrinsic value. This form of the residual income model does not rely on the no-arbitrage assumption – rather it is simply based on the definition of the expected rate of return (the difference between current price and expected cum-dividend end-of-year price divided by current price).

⁹ In Easton, Taylor, Shroff, and Sougiannis (2002) the period t to $t+1$ is 4 years so that eps_{jt+1} is aggregate expected cum-dividend earnings for the four years after date t . We use a one-year forecast horizon instead of four years in order to facilitate more effective use of the data on analysts' recommendations. Easton, Taylor, Shroff, and Sougiannis (2002) note that estimates of the expected rate of return based on just one year of forecasts are very similar to those based on four years of forecasts.

the estimate implied by analysts' earnings forecasts from Easton, Taylor, Shroff, and Sougiannis (2002).

The method adapted from O'Hanlon and Steele (2000) is based on the following form of the residual income valuation model:

$$p_{jt} \equiv bps_{jt} + \frac{(eps_{jt} - r_j \times bps_{jt-1})(1 + g'_j)}{(r_j - g'_j)} \quad (3)$$

The difference between this form of the model and the form used by Easton, Taylor, Shroff, and Sougiannis (2002) is that g'_j is the perpetual growth rate starting from *current residual income* (that is, at time t) that implies a residual income stream such that the present value of this stream is equal to the difference between price and book value; in Easton, Taylor, Shroff, and Sougiannis (2002), g_j is the perpetual growth rate starting from *next-period residual income* (that is, time $t+1$). Since eps_{jt} (that is, realized earnings) is the only pay-off used in estimating the implied expected rate of return based on equation (3), this estimate is not affected by analysts' optimism unless that optimism is shared by the market and captured in p_{jt} .¹⁰ Therefore, the estimate based on current accounting data can serve as an estimate of market expectations. It follows that the difference between the estimate of the expected rate of return based on analysts' forecasts in equation (2) and the estimate based on current earnings in equation (3) is an ex ante estimate of bias introduced when analysts' forecasts are used to estimate the markets' expected rate of return.

2.2. Ex post determination of the effect of bias

Optimistic bias in analysts' earnings forecasts is well-established in the literature; see, for example, O'Brien (1988); Mendenhall (1991); Brown (1993); Dugar and Nathan (1995); and

¹⁰ Our empirical evidence is consistent with the maintained hypothesis that the analysts' optimism is not shared by the market.

Das, Levine, and Sivaramakrishnan (1998). Each of these studies estimates the ex post bias by comparing earnings forecasts with realizations of these forecasted earnings. We obtain an ex post measure of the bias in the estimate of the expected rate of return by comparing the estimate of the expected rate of return based on I/B/E/S analysts' forecasts using the method in Easton, Taylor, Shroff, and Sougiannis (2002) with the expected rate of return based on (perfect foresight forecasts of) earnings realizations; that is, we replace eps_{jt+1}^{IBES} in equation (2) with earnings realizations for period $t+1$, denoted eps_{jt+1}^{PF} . Of course, this ex post comparison, like the studies of bias in analysts' forecasts, will be affected by events having an effect on earnings, which happen between the time of the forecast and the date of the earnings announcement.

2.3. Ex ante and ex post comparisons

In the ex post comparison of expected rates of return, unforeseen events are *omitted* from the market price, which is used as the basis for estimating the expected rate of return. On the other hand, in the ex ante comparison, expectations of future events impounded in market expectations of earnings are not included in the current accounting earnings but are implicitly *included* in the market price, which is used as the basis for estimating the expected rate of return. Since there is no obvious reason to expect a correlation between the information omitted from price in the analyses based on equation (2) and the information included in price but excluded from earnings in the analyses based on equation (3), we use the results from both methods to gain alternative, independent estimates of the bias. As expected our results are similar using either method.

Our maintained hypothesis in the ex ante comparison of implied expected rates of return is that the market at time t sees through (un-does) the optimistic bias in the analysts' forecasts.

The observation that the implied expected rates of return based on current earnings and on realized future earnings are the same, suggests that this maintained hypothesis is reasonable.

2.4. Estimation based on prices, book value, and earnings forecasts

Easton, Taylor, Shroff, and Sougiannis (2002) transform equation (2) to form the following regression relation:

$$\frac{eps_{jt+1}}{bps_{jt}} = \gamma_0 + \gamma_1 \frac{p_{jt}}{bps_{jt}} + \mu_{jt} \quad (4)$$

where $\gamma_0 = g$, $\gamma_1 = r - g$.¹¹ This regression may be estimated for any group/portfolio of stocks to obtain an estimate of the implied expected rate of return, r , and the implied expected growth rate, g , for the portfolio. Easton, Taylor, Shroff, and Sougiannis (2002) run this regression for a sample of U.S. stocks to obtain an estimate of the expected rate of return on the U.S. equity market and hence an estimate of the equity premium for that market. In the empirical implementation of this model, eps_{jt+1} is the I/B/E/S forecast of earnings. Since this is the only pay-off which is used in the estimation of implied expected rate of return, any bias in the forecast will lead to a bias in the estimate of the expected rate of return.

¹¹ At the firm-specific level, the following relation between the regression variables: $\frac{eps_{jt+1}}{bps_{jt}} = \gamma_{0j} + \gamma_{1j} \frac{p_{jt}}{bps_{jt}}$, is

readily obtained by rearranging the identity shown in equation (2). In the re-expression of this relation for a group of observations (as in equation (4)) as a regression relation, the coefficients γ_0 and γ_1 represent an average of the firm-specific γ_{0j} and γ_{1j} coefficients and the cross-sectional variation in these coefficients creates the regression residual. Easton, Taylor, Shroff, and Sougiannis (2002) describe this regression in more detail pointing out that it involves the implicit assumption that it has the properties of a random coefficient regression. It is, of course, possible that the γ_{0j} and γ_{1j} are correlated in cross-section with either (or both) the dependent or the independent variable and this correlation may introduce bias into the estimates of the regression coefficients (and, hence, into the estimates of the implied expected rates of return). It seems reasonable to assume, however, that this bias will be very similar for the regressions based on analysts' earnings forecasts (eps_{jt+1}^{IBES}) and for those based on perfect foresight forecast of earnings (eps_{jt+1}^{PF}). Also, we can think of no reason why the effect of the bias in the analyses based regression (4) will be the same as the effect for the analyses based on current accounting earnings (regression (5)). In other words, similar results from the analysis based on perfect foresight forecasts and from the analyses based on current accounting data support the conclusion that this bias does not unduly affect our estimates.

2.5. Estimation based on current accounting data

The analyses in O'Hanlon and Steele (2000) are based on realized earnings rather than earnings forecasts. Following the essence of the idea in O'Hanlon and Steele (2000), which is summarized in equation (3), we transform this equation to form the following regression relation:¹²

$$\frac{eps_{jt}}{bps_{jt-1}} = \delta_0 + \delta_1 \frac{p_{jt} - bps_{jt}}{bps_{jt-1}} + \zeta_{jt} \quad (5)$$

where $\delta_0 = r$, $\delta_1 = (r - g')/(1 + g')$. This regression may be estimated for any group/portfolio of stocks to obtain an estimate of the expected rate of return, r , and the expected growth rate, g' , for the portfolio. O'Hanlon and Steele (2000) run a regression similar to (5) for a sample of U.K. stocks to obtain an estimate of the expected rate of return on the U.K. equity market; and hence an estimate of the equity premium for that market. In the empirical implementation of regression (5), eps_{jt} is realized earnings. Since this is the only pay-off used in estimating the implied expected rate of return, this estimate is not affected by analysts' optimism unless that optimism is shared by the market and captured in p_{jt} . It follows that the difference between the estimate of the expected rate of return obtained via regression (4) and the estimate based on regression (5) is an ex ante estimate of the bias when analysts' forecasts are used to estimate expected rates of return.

¹² We attribute this model to O'Hanlon and Steele (2000) because they capture its essential elements. The similarity to their model may not, however, be immediately apparent. Since the derivation in O'Hanlon and Steele (2000) is based on Ohlson (1989), the observation that the regression intercept is an estimate of the implied expected rate of return is not evident and O'Hanlon and Steele (2000) do not use it in this way. Rather, they estimate the implied expected rate of return at the firm-specific level by applying their model to time-series data and then measuring the risk premium as the slope of the Securities Market Line estimated from a regression of these firm-specific rates of return on corresponding beta estimates. Notice that, in addition to requiring earnings to be clean surplus in all future periods, this form of the residual income model also requires that the relation between earnings for period t and book value for periods t and $t-1$ follows the clean surplus relation.

2.6. The relation between prices, actual earnings, and forecasts of earnings

In order to ensure that we obtain an estimate of the expected rate of return implied by analysts' forecasts we must use prices in regression (4) that reflect analysts' forecasts. Similarly, in regression (5) we must use prices that reflect earnings realizations to obtain an estimate of the markets' expected rate of return. The alignment of price-dates, earnings announcement dates, and analysts' forecast-dates is described in this sub-section and summarized in figure 1.

We choose the first consensus forecast announced at least 14 days after the date of the earnings announcement.¹³ In the analyses based on these forecasts, we use the price at the close of trade one day after the earnings announcement. Consistent with numerous studies of the information content of earnings, it seems reasonable to assume that this price incorporates the information in realized earnings. Further, we implicitly assume that this price was known to analysts at the time they formed their earnings forecasts. In view of the fact that the forecasts comprising the consensus are formed at various points in time, this assumption may be invalid; some of the forecasts comprising the consensus may precede the earnings announcement date or they may have been issued a considerable time after this date. We examine the sensitivity of the results to this assumption by varying the price-date from the day after the earnings announcement to one day after the consensus forecast is measured. This latter measurement date for price allows for the incorporation of the information in the analysts' forecasts in price. The results are not sensitive to this choice. We will return to this point.

The residual income valuation model underlying regressions (4) and (5) describes the value of a stock at the fiscal period end-date. Our analyses are based on prices after this date. To accommodate this difference, we replace price (p_{jt}) in equations (4) and (5) with price at the

¹³ Use of the first forecast made after the earnings announcement from the I/B/E/S Detail History database does not alter any results.

dates described above discounted by the expected rate of return ($\hat{r}$) back to the fiscal year end; that is, $p_{jt+\tau} / (1 + \hat{r})^{\tau/365}$, where τ is the number of days between the fiscal year end and the price-date. Since the discounting of price requires the expected rate of return we are attempting to estimate in equations (4) and (5), we use an iterative method as used in Easton, Taylor, Shroff, and Sougiannis (2002). We begin these iterations by assuming a discount rate for prices of 12 percent. We run each regression and obtain estimates of the expected rate of return which we then use as the new rate for discounting prices. We then re-run the regressions to re-estimate equation (4) and/or equation (5) and provide another estimate of expected return. This procedure is repeated until the estimate of the expected return and the rate used in discounting price converge.¹⁴

3. Description of the data

All earnings forecast and recommendation data are obtained from the I/B/E/S unadjusted research databases. We use the first median consensus forecast of earnings for year $t+1$ released 14 days or more after the announcement of earnings for year t . This forecast is released on the third Thursday of each month. These data are obtained from the I/B/E/S Summary database. “Actual” earnings are also obtained from this database. The first year of our analyses uses forecasts and recommendations for 1993 in order to ensure the dates of the individual analysts’ forecasts are reliable.¹⁵ Book value of common equity and common shares outstanding are

¹⁴ This iterative process is repeated until none of the annual estimates changes by more than 0.00001%. In our samples, the annual estimates usually converged in 5-6 iterations. This iterative procedure is not sensitive to choices of beginning discount rates between five and 20 percent.

¹⁵ Zitzewitz [2002, p. 16] describes the importance of not relying on forecast dates in the I/B/E/S database prior to 1993 as follows:

“I/B/E/S dates forecasts using the date it was entered into the I/B/E/S system. It has been well documented (e.g., by O’Brien, 1988) that the lags between a forecast becoming public and its entry into the I/B/E/S system were substantial in the 1980s (i.e., up to a month). In the 1980s, analysts mailed their forecasts,

obtained from the CRSP/COMPUSTAT annual merged database.¹⁶ Prices are obtained from the CRSP daily price file.

We delete firms with non-December fiscal-year end so that the market implied discount rate and growth rate are estimated at the same point in time for each firm-year observation. For each set of tests, firms with any of the dependent or independent variables for that year in the top or bottom two percent of observations are removed to reduce the effects of outliers. Dropping between one and five percent of observations does not affect the conclusions of the study. For December 1999, in particular, removal of only one percent of observations has a large effect on that year's results in the value-weighted analyses; this is due to the extremely high price-to-book ratios of some internet firms prior to the market crash in 2000.

4. Ex post and ex ante bias in analysts' consensus forecasts

We begin by documenting the accuracy (that is, the mean/median *absolute* earnings forecast error) and the ex post bias (that is, the mean/median earnings forecast error) in the earnings forecasts for the entire sample of stocks. We then compare the estimate of the expected rate of return implied by prices, book values, and analysts' forecasts of earnings with the

often in monthly batches, to I/B/E/S where they were hand entered into the system. Since 1991-92, however, almost all analysts have entered their forecasts directly into the I/B/E/S system on the day they wish to make their forecast widely available (Kutsoati and Bernhardt, 1999). Current practice for analysts is now usually to publicly release forecasts within 24 hours of providing them to clients. I/B/E/S analysts have real-time access to each other's forecasts through this system, so an analyst entering a forecast into the system on Wednesday knows about forecasts entered on Tuesday and could potentially revise her forecast to incorporate their information. An additional advantage of the post-92 data is the shift from retrospective data entry by a specialist to real-time data entry by either the analyst or her employee should have considerably reduced data-entry related measurement error."

¹⁶ In order to ensure that the clean-surplus assumption required for the derivation of the residual income valuation model holds in the data for fiscal year t , contemporaneous book value in regression (5) – that is, b_{jt} – is calculated as Compustat book value of common equity minus Compustat net income plus I/B/E/S actual income. That is, we use the book value number that would have been reported if the (corresponding) income statement had been based on I/B/E/S actual earnings. We also remove year t dirty surplus items from Compustat book value. These adjustments are unnecessary for the book value variable in regression (4) because the clean-surplus assumption only refers to future income statements and balance sheets.

estimate obtained from prices, book values, and actual current earnings. This is an estimate of ex ante bias in the estimates of the expected rate of return reported in the extant literature.

4.1. Accuracy and bias in the analysts' forecasts of earnings

Table 1 summarizes the accuracy and the ex post measure of bias in the I/B/E/S consensus forecast of earnings at the end of each of the years 1992 to 2003. We use the mean and the median absolute forecast error as the measure of accuracy. The mean absolute forecast error ranges from \$0.427 in 1994 to \$1.394 in 2000; the median absolute forecast error ranges from \$0.160 in 2002 to \$0.310 in 2000. We also present the mean and the median absolute forecast error deflated by end-of-year price in order to give an indication of the scale of these errors. The mean absolute price-deflated forecast error ranges from 0.019 in 2003 to 0.052 in 2000; the median absolute price-deflated forecast error ranges from 0.008 in 2003 to 0.018 in 2000.

We use the mean (median) forecast error as the measure of the ex post bias in the analysts' forecasts. The mean forecast error ranges from -\$1.257 in 2000 to \$0.119 in 2002. The median forecast error ranges from -\$0.240 in 2000 to -\$0.010 in 2003. The mean price-deflated forecast error ranges from -0.041 in 2000 to -0.003 in 2003. The median price-deflated forecast error ranges from -0.012 in 2000 to 0.000 in 2003.

These predominantly negative forecast errors are consistent with the prior literature, which concludes that analysts' forecasts, particularly long-run forecasts, tend to be optimistic; see, for example, O'Brien (1993); Lin (1994); and Richardson, Teoh, and Wysocki (2004). As noted earlier, these forecast errors compare forecasts with ex post realizations.

4.2 Description of regression variables

The number of observations we use to estimate the annual regressions ranges from 1,418 at December 1992 to 2,137 at December 1997. As shown in table 2, the mean price-to-book ratio, which is the independent variable in regression (4), ranges from 1.945 at December 2002 to 3.398 at December 1999; the median price-to-book ratio ranges from 1.625 at December 2002 to 2.409 at December 1997. Regression (4) is run with the forecasted return-on-equity based on the I/B/E/S consensus forecast as the dependent variable. The mean forecasted return-on-equity ranges from 0.079 at December 2001 to 0.146 at December 1994; the median forecasted return-on-equity ranges from 0.111 at December 2001 to 0.145 at December 1994.

The annual mean and median current return-on-equity, which is the dependent variable in regression (5), is generally a little less than the corresponding mean and median forecasted return-on-equity. The mean current return-on-equity ranges from 0.077 at December 2001 to 0.122 at December 1995; the median current return-on-equity ranges from 0.010 at December 2001 to 0.132 at December 1995. The mean of the independent variable in this regression, the difference between price and current book value deflated by lagged book value, ranges from 1.007 at December 2002 to 2.699 at December 1999; the median ranges from 0.662 at December 2002 to 1.491 at December 1997.

4.3. Comparison of implied expected rates of return based on I/B/E/S forecasts of earnings with implied expected rate of return based on current accounting data

In this section, we compare the estimates of the implied expected rates of return based on the method in Easton, Taylor, Shroff, and Sougiannis (2002), which uses one-year ahead I/B/E/S consensus forecasts of earnings in regression (4), with the estimates obtained from the method adapted from O'Hanlon and Steele (2000), which uses current earnings and current and lagged

book value in regression (5). We also compare the estimates based on analysts' forecasts to those implied by future earnings realizations; that is, by perfect foresight forecasts.

4.3.1. The expected rate of return implied by analysts' earnings forecasts

The summary statistics from regression (4), where the dependent variable is I/B/E/S forecasted return-on-equity, are included in panel A of table 3. We provide year-by-year estimates of the regression coefficients and t-statistics for tests of their difference from zero. These t-statistics may be over-stated due to the possibility of correlated residuals; so we present the mean coefficient estimates and the related Fama and MacBeth (1973) t-statistics. The regression adjusted r-square ranges from 0.73 percent at December 1999, to 36.60 percent at December 1992.¹⁷ The mean estimate of the intercept coefficient γ_0 , an estimate of the implied growth in residual income beyond the one-year forecast horizon, is 0.074 with a t-statistic of 8.50. The mean estimate of the slope coefficient γ_1 , an estimate of the difference between the implied expected rate of return and the implied growth in residual income beyond the one-year forecast horizon, is 0.020 with a t-statistic of 5.86.

The estimates of the implied expected rate of return obtained from the estimates of the regression (4) coefficients, where the dependent variable is analysts' forecasts of return-on-equity, are in panel A of table 3. These estimates range from 4.93 percent at December 2001, to 13.29 percent at December 1999; with a mean (t-statistic) of 9.43 percent (14.16).

¹⁷ We note the very low r-square in some of these regressions. As a result we performed several analyses of the effects of outliers including more severe outlier removal – for example, removing up to the top and bottom 20 percent of observations or by eliminating all observations with an R-student statistic greater than 2 -- the regression r-square increases but none of our inferences based on the resulting estimates of the implied expected rate of return change. We also perform all analyses on the sub-set of observations for which analysts forecast positive earnings. Again we obtain much higher r-squares but inferences remain unchanged. These further analyses of outliers are also performed on all subsequent regressions and, in all cases, our inferences are unchanged.

4.3.2. The expected rate of return implied by current accounting data

The summary statistics from regression (5) are included in panel A of table 3. The regression adjusted r-square ranges from 0.34 percent at December 1999 to 27.09 percent at December 1992. The mean estimate of the intercept coefficient δ_0 , which is an estimate of the implied expected rate of return, is 0.066 (t-statistic of 10.50); and the mean estimate of the slope coefficient δ_1 , which is a function of the expected rate of return and the expected growth in residual income, is 0.022 (t-statistic of 5.51). The estimates of the implied expected rate of return are also included in panel A of table 3. These estimates range from 2.82 percent at December 2001 to 9.97 percent at December 1999; with a mean (t-statistic) of 6.59 percent (10.50).

4.3.3. The ex ante difference between the estimate of the expected rate of return based on analysts' earnings forecasts and the estimate of the expected rate of return based on current accounting data

Differences between the estimates of expected rate of return based on regressions (4) and (5) are included in the last column of panel A of table 3. On average, the difference between the estimate of the expected rate of return based on analysts' earnings forecasts and the estimate of the expected rate of return based on earnings realizations is 2.84 percent (t-statistic of 12.33). There are some years when the difference is quite large; for example, for the sample of stocks at December 1994, the difference is 3.83 percent. These results are not surprising in view of the fact that analysts' forecasts are known to be optimistic.

An implication of the observation that expected rates of return based on analysts' forecasts tend to be higher is that caution should be taken when interpreting the meaning of the rate of return that is implied by analysts' earnings forecasts; if, as is often the case in the extant literature, it is used as an estimate of the cost of capital, it is likely upward biased.

4.3.4. Estimates of the expected rate of return based on perfect foresight forecasts

The results in section 4.3.3 are roughly consistent with the results in Table 1. For example, we saw, in Table 1 that the mean deflated forecast error is -0.020. A crude PE valuation model, which relies on full payout and earnings following a random walk, suggests that the price-to-forward-earnings ratio is equal to the inverse of the expected rate of return. Thus a deflated forecast error of -0.020 implies an error in the expected rate of return of 2 percent. Allowing for the conservative nature of accounting, as in the models used in the ex ante indicators of optimism in panel A of table 3, leads to the conclusion that these estimates are at least “in the same ball-park”.

Alternatively, the ex post forecast error can be re-parameterized as an error in the implied expected rate of return. This error may be estimated as the difference between the implied expected rate of return based on regression (4) where expected earnings are I/B/E/S forecasts (as in panel A of table 3) and the implied expected rate of return when these expected earnings are replaced in this regression with realized earnings for year $t+1$. The results of estimating the implied expected rate of return using realized earnings as “perfect foresight” forecasts are reported in panel B of table 3. Using perfect foresight earnings, the estimates of expected rate of return range from 3.13 percent at December 2001 to 9.79 percent at December 1999; with a mean (t-statistic) of 6.68 percent (10.79). Comparing the perfect foresight forecast to the consensus forecasts, the mean bias is 2.75 percent (t-statistic of 7.13).

4.3.5. Comparison of the estimates of the expected rate of return

The two estimates of expected rate of return that are not expected to contain bias, that is, those based on perfect foresight earnings and those based on current accounting data are very similar. The difference of -0.09 percent between these estimates is not significantly different

from zero with a t -statistic of -0.19. It follows that our estimates of the bias are similar using either method. That is, both methods yield alternative, independent estimates of the bias that do not differ significantly; this observation supports the maintained hypothesis that the market sees through the optimistic bias in the analysts' forecasts.

Further evidence consistent with the notion that the market sees through the optimistic bias is the fact that, consistent with Richardson, Teoh, and Wysocki (2004), the forecast error declines almost monotonically as the forecast horizon decreases from approximately 12 months as in the analyses in panel C of table 3 to shortly before the earnings announcement date for year $t+1$. The un-tabulated associated implied expected rate of return based on these forecast and prices immediately following these forecasts also decreases almost monotonically to 6.47 percent for the consensus forecasts (of $t+1$ earnings) made in January of year $t+1$. That is, the expected rate of return implied by analysts' forecasts declines to the expected rate of return implied by the ex ante estimate of the expected rate of return implied by accounting earnings at date t . Again these results suggest that the market at date t sees through the optimistic bias in the analysts' forecasts of earnings for period $t+1$.

4.3.6. Effects of altering the timing of price measurement

As mentioned in section 2.3, we use price measured after the release of the prior year earnings but before analysts' forecast revisions in our primary analyses. Panel C of table 3 summarizes the results of the analysis summarized in panels A and B of table 3, but using prices measured at close of trade on the day after the consensus forecast is measured. This price is at least 14 days and could be a month and a half after the price used in panels A and B. We assume that this price reflects the information in the analysts' forecasts. Comparison of panels A and C reveals that the measurement of price at differing points; and, therefore, differing periods for

discounting of price back to fiscal year-end; has no statistically or economically significant effect. The primary result from panel A of table 3 of an average 2.84 percent difference between the analysts' and market's expected rate of return is virtually unchanged at 2.93, with an untabulated t-statistic of 14.69, when price is measured at the day after the consensus forecast is measured.¹⁸

5. Value-weighted estimates of the implied expected rate of return

The analyses in section 4 examine the average effect of bias in analysts' forecasts of earnings on estimates of the implied expected rate of return. All observations are given equal weight in the analyses. Such weighting will be appropriate in some studies. Easton, Sommers, and Zmijewski (2006), for example, compare the difference between the expected rate of return implied by analysts' forecasts and the expected rate of return implied by current earnings for firms subject to litigation under section 10b-5.¹⁹ Since the focus of their study is on average differences, they give each observation equal weight; value-weighting would lead to results that were dominated by cases associated with WorldCom and Enron.

Value-weighting will be more appropriate in many studies. Perhaps the best example is the estimation of the equity risk premium, which is a central part of three well-known studies based on analysts' earnings forecasts by Gebhardt, Lee, and Swaminathan (2001); Claus and Thomas (2001); and Easton, Taylor, Shroff, and Sougiannis (2002). These studies give equal weighting to all stocks. Yet, estimating the risk premium from investing in the equity market is more meaningful if stocks are weighted by their market capitalization. In the equally-weighted

¹⁸ The results are virtually identical if we use prices taken from any date ranging from one day after the earnings announcement date to one day after the forecast announcement date (the set of s price-dates shown in Figure 1).

¹⁹ Under Rule 10b-5, a firm and its officials can be held liable for damages to investors who bought and sold the firm's securities if the damages are attributable to investors' reliance on misleading statements or omission of material facts.

analyses in the papers referred to above, small stocks will have an undue effect on the estimate of the market return. Further, stocks with low or negative earnings, which are somewhat meaningless as summary valuation metrics, potentially have an influence that is similar to the influence of large stable firms where earnings are a much more meaningful valuation metric. In order to avoid these undue influences, and to provide an estimate of the equity risk premium that is (1) not affected by analysts' optimism; and (2) more representative of the risk premium for the market portfolio; we repeat all of the analyses weighting each of the observations by market capitalization.

In order to provide a sense of the likely effect of value weighting, we begin by describing the way that analysts' optimism differs with firm size. We also document the relation between firm size and the variables used in regressions (4) and (5). Central to our analyses is the observation, documented in panel A of table 4, that the mean scaled absolute forecast error declines in a monotonic manner from 0.102 for the decile of smallest firms to 0.012 for the decile of largest firms. Similarly, the median absolute scaled forecast error declines in a monotonic manner from 0.042 to 0.006.

Analysts' optimism, measured by the mean (median) forecast error, declines almost monotonically from -0.116 (-0.023) for the decile of smallest firms to -0.086 (-0.002) for the decile of largest firms. The differences in optimistic bias across these size deciles illustrate the point that difference in bias across samples of observations may explain a significant portion of the difference in the implied expected rates of return across these samples; in other words, differences in bias across samples may lead to spurious inferences.

Consistent with prior literature, see, for example, Fama and French (1992), the price-to-book ratio increases with firm size from a mean of 1.707 for the decile of smallest firms to a

mean of 3.593 for the decile of largest firms. The forecasted and the realized return-on-equity also increase with firm size, suggesting that the smaller firms tend to be firms with higher expected earnings growth.²⁰

The results from the estimation of value-weighted regressions (4) and (5) are summarized in panel B of table 4. A notable difference between these value-weighted regression results and the results for equally-weighted regressions (see panels A and B of table 3) is the higher adjusted r-square for the value-weighted regressions. For example, the average adjusted r-square for regression (4) based on analysts' consensus forecasts is 47.16 percent for the value-weighted regression; whereas it is 9.58 percent for the equally-weighted regression. As expected, t-statistics on the coefficient estimates in these value-weighted regressions are also higher.

The mean estimates (t-statistic) of the expected rate of return, also reported in panel B of table 4, are 11.27 percent (21.20) using analysts' forecasts and 9.67 percent (13.90) using current accounting data.²¹ The un-tabulated minimum expected rate of return estimated using current accounting data is 6.22 percent at December 1992. The average of 9.67 percent yields a more reasonable estimate of the risk premium than the equal-weighted sample; 4.43 percent using 5-year treasuries as a proxy for the risk free rate. Differences between the estimates are also reported in panel B of table 4. The difference, though smaller in the value-weighted analyses than in the equally-weighted analyses, 1.60 percent compared with 2.84 percent, is still significantly positive (t-statistic of 4.90).

²⁰ The firms in the deciles of smaller firms also tend to have a much greater proportion of losses (the proportion of losses decreases monotonically from 17.64 percent for the decile of smallest firms to 1.65 percent for the decile of largest firms).

²¹ The mean estimate (t-statistic) of the expected rate of return based on perfect foresight forecasts is 10.63 percent (14.35).

6. Variation in the implied expected rate of return with changes in the percentage of analysts making “buy” recommendations

Having documented a bias in the estimates of the expected rate of return based on analysts’ forecasts of earnings, we now examine how the bias varies across analysts’ recommendations. It is well-known that analysts seldom issue “sell” recommendations. To the extent that our samples examined thus far contain a majority of firms with “buy” recommendations, the observed positive bias in the expected rate of return using analysts’ forecasts may be capturing the analysts’ expectation of the abnormal returns, which can be earned from these stocks. To examine this notion, we compare estimates of the expected rates of return for stocks where the consensus forecast is comprised of analysts with varying recommendation types.

6.1 Sample description

I/B/E/S provides data on the percentage of analysts whose forecasts comprise the consensus who also make either a “strong buy” or a “buy” recommendation. We repeat the analyses in section 4.3 for sub-samples with various percentages of these types of recommendations. Descriptive statistics are provided in table 5, panel A. The choice of the five partitions of the data is based on a desire to maintain a sufficient number of observations to provide reasonable confidence in the regression output in each year. We restrict the sample to those consensus forecasts which are comprised of at least 5 analysts so that it is possible for a firm to appear in any of the partitions.²²

The mean and median forecast error is always negative; that is, analysts are optimistic, regardless of the percentage of “buy” recommendations in the consensus. For example, the median deflated forecast error is -0.004 when the percentage of buy recommendations is greater

²² Our findings and conclusions are unchanged when firms with consensus forecasts comprised of less than 5 analysts are included.

than 90 percent, between 30 and 50 percent, and when the percentage of “buy” recommendations is less than 30 percent.

Both the return-on-equity and the price-to-book ratio tend to be higher for the observations where there are more “buy” recommendations comprising the consensus. For example, the median forecasted return-on-equity for the sub-samples where greater than 90 percent of the analysts recommend “buy” and where between 70 and 90 percent recommend “buy” is 0.157 and 0.162 while median forecasted return-on-equity for the sub-sample where less than 30 percent of the analysts recommend “buy” is 0.112. The median price-to-book ratio for the sub-samples where greater than 90 percent of the analysts recommend “buy” and where between 70 and 90 percent recommend “buy” is 3.011 and 2.686 while median price-to-book ratio for the sub-samples where less than 30 percent of the analysts recommend “buy” is 1.649.

6.2. Estimates of implied expected rates of return

The results from the estimation of regression (4) based on price, I/B/E/S forecasts of earnings, and current book value and from the estimation of regression (5) based on price and current accounting data and are summarized in table 5, panel B. We focus our discussion on the estimates of the implied expected rates of return obtained from these regression parameters. These estimates are also included in panel B.

The estimates of the expected rates of return implied by I/B/E/S analysts’ forecasts decline almost monotonically with the percentage of “buy” recommendations associated with the forecasts of earnings comprising the consensus; the means of these estimates are 11.20 percent, 11.84 percent, 10.82 percent, 9.18 percent, and 6.86 percent, suggesting that analysts’ recommendations are, indeed, consistent with the implied expectations of rates of return. The estimates of the expected rates of return based on prices and current accounting data show a

pattern that is very similar to that of those based on analysts' forecasts. The mean estimates of the expected rate of return for each of the groups of data decline monotonically with the percentage of "buy" recommendations associated with the forecasts of earnings comprising the consensus; the means of these estimates are 10.94 percent, 10.22 percent, 8.90 percent, 7.23 percent, and 4.60 percent.

Differences between the estimates of expected rate of return based on percentage of "buy" recommendations are included in table 5, panel C. Comparing the expected rates of return based on prices and current accounting data with the estimates based on analysts' forecasts reveals that even when the analysts are not recommending "buy" their forecasts imply a rate of return that is higher than expectations based on current accounting data; these mean differences between the estimates based on analysts' forecasts and estimates based on current accounting data are 0.26 percent, 1.61 percent, 1.92 percent, 1.95 percent, and 2.27 percent. Four of these differences are significant. This pervasive optimism in the expected return measured by comparing analysts' return expectations with return expectations based on current accounting data is, interestingly, quite similar to the pervasive optimism observed when comparing expectations of future earnings with actual realizations of earnings; see table 5, panel A.

6.3. Summary

To summarize the analyses in this section, we observe that analysts' recommendations are consistent with their expectations of returns; that is, there is a monotonic decrease in expected rate of return as the percentage of "buy" recommendations declines.²³ Analysts' expected rates of return are higher than expectations based on current accounting data regardless of their recommendation. An interpretation of this result is that analysts are always optimistic;

²³ Our findings and conclusions are unchanged when the analysis is repeated using a value-weighted analysis similar to section 5.

even when they are not issuing “buy” recommendations.²⁴ The bias in expected rates of return based on analysts’ forecasts is not the result of analysts’ expectations of positive abnormal returns isolated in firms with “buy” or “strong buy” recommendations.

7. Summary and conclusions

We show that, on average, the difference between the estimate of the expected rate of return based on analysts’ earnings forecasts and the estimate of based on current earnings realizations is 2.84 percent. An implication of the observation that rates of return based on analysts’ forecasts are higher than market expectations is that caution should be taken when interpreting the meaning of the rate of return that is implied by analysts’ earnings forecasts; it may not be, as the literature generally claims, an estimate of the cost of capital.

When estimates of the expected rate of return in the extant literature are adjusted to remove the effect of optimism bias in analysts’ forecasts, the estimate of the equity risk premium appears to be approximately zero. We show, however, when estimates are based on value-weighted analyses, the bias in the estimate of the expected rate of return is lower and the estimate of the expected equity premium is more reasonable; 4.43 percent.

Results from sub-samples formed on the basis of percentage of analysts comprising the consensus recommending “buy” show that the estimate of the expected rate of return, based on both analysts’ forecasts of earnings and on current earnings, declines in a monotonic manner as the percentage of analysts recommending “buy” declines. A comparison of the estimates of the expected rate of return based on the analysts’ forecasts, with estimates based on earnings realizations, suggests that analysts tend to be more optimistic than the market even when they are

²⁴ This result is consistent with Barber, Lehavy, McNicholls, and Trueman (2001) who show that analysts’ recommendations (in their case, those summarized in the Zach’s database) can not be used to form profitable trading strategies.

not making “buy” recommendations. That is, analysts recommend “buy” when they expect the future return to be high and “sell” when they expect the return to be low regardless of market expectations.

Our paper has two key implications for future research which uses market price, book value of equity, and accounting earnings to obtain estimates of the implied expected rate of return for a portfolio of stocks. First, since analysts’ forecasts are pervasively optimistic, estimates of the implied expected rate of return formed using forecasts will be pervasively and significantly upward biased. This bias may be avoided by estimating the rate of return implied by price, book values, and *realized* earnings rather than biased earnings *forecasts*. Second, value-weighted analyses may be more appropriate in addressing certain issues such as estimating the equity premium, than equal-weighted analyses. The value-weighted analyses may provide more realistic estimates of the expected rate of return than are implied by equally-weighted analyses; which may be unnecessarily affected by less representative observations, such as penny stocks, and stocks making losses.

When coupled with results from the papers that demonstrate the troublesome effects of measurement error in firm-specific estimates of the expected rate of return, the results in this study suggest that the extant measures of implied expected rate of return should be used with considerable caution. The challenge is to find means of reducing the measurement error and to mitigate the effects of bias. Easton and Monahan (2005) suggest focusing on sub-samples where the measurement error is likely to be small. Our paper suggests that methods based on realized earnings rather than earnings forecasts may be a possible means of avoiding the effects of bias in analysts’ forecasts. Another possible avenue might be to attempt to un-do the bias; following, for example, the ideas in Frankel and Lee (1998).

Figure 1: Alignment of Price-Dates, Earnings Announcement Dates, and Analysts' Forecast-Dates




Table 1: Descriptive statistics on forecast errors for the consensus sample

t	N	Accuracy of forecasts				Bias in forecasts			
		$ FE_{jt+1} $		$ FE_{jt+1} /p_{jt}$		FE_{jt+1}		FE_{jt+1}/p_{jt}	
		Mean	Median	Mean	Median	Mean	Median	Mean	Median
12/92	1,418	0.594	0.280	0.030	0.014	-0.241	-0.150	-0.017	-0.007
12/93	1,544	0.461	0.190	0.028	0.009	-0.228	-0.070	-0.019	-0.003
12/94	1,781	0.427	0.220	0.030	0.012	-0.206	-0.080	-0.019	-0.004
12/95	1,939	0.451	0.210	0.028	0.011	-0.261	-0.070	-0.019	-0.004
12/96	2,006	0.518	0.210	0.027	0.010	-0.187	-0.100	-0.018	-0.005
12/97	2,137	0.606	0.270	0.031	0.013	-0.376	-0.200	-0.024	-0.009
12/98	2,044	0.718	0.215	0.040	0.012	-0.515	-0.080	-0.025	-0.004
12/99	1,854	0.668	0.230	0.046	0.012	-0.399	-0.090	-0.028	-0.004
12/00	1,729	1.394	0.310	0.052	0.018	-1.257	-0.240	-0.041	-0.012
12/01	1,809	0.705	0.200	0.033	0.011	0.063	-0.060	-0.018	-0.003
12/02	1,825	0.570	0.160	0.031	0.011	0.119	-0.030	-0.012	-0.002
12/03	2,000	0.650	0.170	0.019	0.008	-0.251	-0.010	-0.003	0.000
Means	1,841	0.647	0.222	0.033	0.012	-0.312	-0.098	-0.020	-0.005

Notes to Table 1:

FE_{jt+1} is actual earnings per share for year $t+1$ as reported by I/B/E/S less the first median consensus forecast of earnings per share for year $t+1$ released at least 14 days after the announcement of year t earnings

p_{jt} is price per share as of the end of fiscal year t

Table 2: Summary statistics for regression variables

<i>t</i>	N	$\frac{eps_{jt+1}^{Cons}}{bps_{jt}}$		$\frac{eps_{jt}}{bps_{jt-1}}$		$\frac{p'_{jt}}{bps_{jt}}$		$\frac{p'_{jt} - bps_{jt}^*}{bps_{jt-1}}$	
		Equation (4) dependent variable		Equation (5) dependent variable		Equation (4) independent variable		Equation (5) independent variable	
		Mean	Median	Mean	Median	Mean	Median	Mean	Median
12/92	1,418	0.138	0.132	0.104	0.110	2.193	1.792	1.265	0.854
12/93	1,544	0.138	0.138	0.113	0.122	2.374	1.929	1.505	0.994
12/94	1,781	0.146	0.145	0.121	0.126	2.114	1.706	1.334	0.834
12/95	1,939	0.145	0.142	0.122	0.132	2.454	1.906	1.679	1.060
12/96	2,006	0.135	0.139	0.108	0.126	2.654	2.114	1.851	1.228
12/97	2,137	0.125	0.140	0.102	0.125	2.998	2.409	2.132	1.491
12/98	2,044	0.118	0.134	0.093	0.116	2.728	1.974	1.810	0.959
12/99	1,854	0.126	0.141	0.094	0.124	3.398	1.883	2.699	0.996
12/00	1,729	0.116	0.136	0.100	0.130	2.749	1.964	2.022	1.109
12/01	1,809	0.079	0.111	0.068	0.100	2.457	1.928	1.548	0.989
12/02	1,825	0.093	0.117	0.077	0.102	1.945	1.625	1.007	0.662
12/03	2,000	0.106	0.121	0.090	0.111	2.883	2.314	2.198	1.450
Means	1,841	0.122	0.133	0.099	0.119	2.579	1.962	1.754	1.052

Notes to Table 2:

eps_{jt+1}^{Cons}	is the first median consensus forecast of earnings per share for firm <i>j</i> for year <i>t</i> +1 released at least 14 days after the announcement of year <i>t</i> earnings
eps_{jt}	is the I/B/E/S actual earnings per share for firm <i>j</i> for year <i>t</i>
bps_{jt}	is common book value of equity per share for firm <i>j</i> at time <i>t</i>
$p'_{jt} = \frac{P_{jt+\tau}}{(1+\hat{r})^{365}}$	is the price per share for firm <i>j</i> at time <i>t</i> + <i>τ</i> (one day after the earnings announcement date), $P_{jt+\tau}$, adjusted for stock splits and stock dividends since the end of the fiscal year, discounted to year end using the estimated discount rate
bps_{jt}^*	is the common book value of equity per share for firm <i>j</i> at time <i>t</i> less net income for firm <i>j</i> for year <i>t</i> plus I/B/E/S actual earnings per share for firm <i>j</i> for year <i>t</i>

Table 3: Comparison of implied expected rates of return based on I/B/E/S forecasts of earnings with implied expected rate of return based on current accounting data

Panel A: Estimates of expected rate of return based on analysts' forecasts and current accounting data

$\frac{eps_{jt+1}^{Cons}}{bps_{jt}} = \gamma_0 + \gamma_1 \frac{p'_{jt}}{bps_{jt}} + \mu_{jt} \quad (4)$						$\frac{eps_{jt}}{bps_{jt-1}} = \delta_0 + \delta_1 \frac{p'_{jt} - bps_{jt}^*}{bps_{jt-1}} + \zeta_{jt} \quad (5)$				Difference in expected rate of return
Analysts' consensus earnings forecasts						Current accounting data				
<i>T</i>	N	γ_0	γ_1	Adj R ²	$\hat{r} = \gamma_0 + \gamma_1$	δ_0	δ_1	Adj R ²	$\hat{r} = \delta_0$	
12/92	1,418	0.057 (17.71)	0.037 (28.62)	36.60%	9.39%	0.057 (18.96)	0.037 (22.97)	27.09%	5.67%	3.72%
12/93	1,544	0.073 (16.53)	0.027 (16.91)	15.59%	10.08%	0.068 (18.37)	0.030 (16.74)	15.32%	6.83%	3.25%
12/94	1,781	0.073 (16.25)	0.035 (18.99)	16.81%	10.73%	0.069 (21.01)	0.039 (23.73)	24.00%	6.90%	3.83%
12/95	1,939	0.095 (23.47)	0.021 (15.38)	10.83%	11.53%	0.092 (23.40)	0.018 (11.70)	6.55%	9.22%	2.31%
12/96	2,006	0.089 (18.91)	0.018 (12.00)	6.66%	10.61%	0.073 (16.79)	0.019 (12.11)	6.77%	7.26%	3.35%
12/97	2,137	0.082 (14.64)	0.014 (9.13)	3.71%	9.64%	0.066 (14.61)	0.017 (11.30)	5.60%	6.62%	3.02%
12/98	2,044	0.082 (15.23)	0.013 (8.67)	3.50%	9.50%	0.065 (15.86)	0.016 (11.89)	6.43%	6.49%	3.01%
12/99	1,854	0.136 (32.67)	-0.003 (-3.83)	0.73%	13.29%	0.100 (22.54)	-0.002 (-2.71)	0.34%	9.97%	3.32%
12/00	1,729	0.084 (15.42)	0.012 (7.84)	3.38%	9.57%	0.086 (16.02)	0.007 (4.30)	1.00%	8.61%	0.96%
12/01	1,809	0.029 (4.64)	0.020 (9.42)	4.63%	4.93%	0.028 (6.30)	0.026 (14.20)	9.99%	2.82%	2.11%
12/02	1,825	0.019 (3.12)	0.038 (14.14)	9.83%	5.70%	0.030 (7.98)	0.047 (22.13)	21.13%	2.96%	2.74%
12/03	2,000	0.069 (11.65)	0.013 (7.55)	2.72%	8.18%	0.057 (11.55)	0.015 (9.59)	4.35%	5.74%	2.44%
Means	1,841	0.074 (8.50)	0.020 (5.86)	9.58%	9.43%	0.066 (10.50)	0.022 (5.51)	10.71%	6.59%	2.84%
t-Statistics					(14.16)				(10.50)	(12.33)

Table 3: Continued

Panel B: Estimates of expected rate of return based on future realized earnings

	$\frac{eps_{jt+1}}{bps_{jt}} = \gamma_0 + \gamma_1 \frac{p'_{jt}}{bps_{jt}} + \mu_{jt} \quad (4)$				Analysts' Forecasts Less Perfect Foresight	Current Accounting Data Less Perfect Foresight
	Perfect foresight earnings forecasts					
<i>t</i>	γ_0	γ_1	Adj R ²	$\hat{r} = \gamma_0 + \gamma_1$		
12/92	0.037 (7.09)	0.031 (15.31)	14.10%	6.77%	2.62%	-1.10%
12/93	0.049 (8.10)	0.026 (11.61)	7.97%	7.45%	2.63%	-0.62%
12/94	0.046 (7.56)	0.031 (12.77)	8.33%	7.71%	3.02%	-0.81%
12/95	0.076 (13.29)	0.013 (6.69)	2.22%	8.87%	2.66%	0.35%
12/96	0.082 (12.01)	0.004 (1.83)	0.12%	8.56%	2.05%	-1.30%
12/97	0.040 (5.14)	0.009 (4.18)	0.77%	4.89%	4.75%	1.73%
12/98	0.057 (8.28)	0.006 (3.15)	0.44%	6.27%	3.23%	0.22%
12/99	0.105 (17.73)	-0.007 (-6.01)	1.87%	9.79%	3.50%	0.18%
12/00	0.043 (6.16)	0.004 (2.05)	0.18%	4.70%	4.87%	3.91%
12/01	0.018 (2.47)	0.013 (5.16)	1.40%	3.13%	1.80%	-0.31%
12/02	-0.003 (-0.48)	0.041 (13.60)	9.16%	3.77%	1.93%	-0.81%
12/03	0.075 (11.02)	0.007 (3.71)	0.64%	8.28%	-0.10%	-2.54%
Means	0.052	0.015	3.93%	6.68%	2.75%	-0.09%
t-Statistics	(6.12)	(3.63)		(10.79)	(7.13)	(-0.19)

Table 3: Continued

Panel C: Comparison of implied expected rates of return based on I/B/E/S forecasts of earnings with implied expected rate of return based on current accounting data and on future realized earnings using prices measured the day after the consensus forecast

$$\frac{eps_{jt+1}^{Cons}}{bps_{jt}} = \gamma_0 + \gamma_1 \frac{p'_{jt}}{bps_{jt}} + \mu_{jt} \quad (4)$$

Analysts' consensus earnings forecasts

	N	γ_0	γ_1	Adj R ²	$\hat{r} = \gamma_0 + \gamma_1$
Means	1,841	0.072	0.021	10.07%	9.34%
t-Statistics		(8.04)	(5.93)		(13.68)

$$\frac{eps_{jt}}{bps_{jt-1}} = \delta_0 + \delta_1 \frac{p'_{jt} - bps_{jt}^*}{bps_{jt-1}} + \zeta_{jt} \quad (5)$$

Current accounting data

	N	δ_0	δ_1	Adj R ²	$\hat{r} = \delta_0$
Means	1,841	0.064	0.023	11.36%	6.41%
t-Statistics		(10.13)	(5.86)		(10.13)

$$\frac{eps_{jt+1}^{PF}}{bps_{jt}} = \gamma_0 + \gamma_1 \frac{p'_{jt}}{bps_{jt}} + \mu_{jt} \quad (4)$$

Perfect foresight earnings forecasts

	N	γ_0	γ_1	Adj R ²	$\hat{r} = \gamma_0 + \gamma_1$
Means	1,841	0.049	0.016	4.42%	6.50%
t-Statistics		(5.36)	(3.84)		(9.72)

Notes to Table 3:

Panel A of the table reports the results of estimating regression (4) using I/B/E/S consensus forecasts and regression (5) using current accounting data cross-sectionally using all available observations. Panel B reports the results of estimating regression (4) using subsequent earnings realizations as perfect foresight forecasts. Observations with any of the dependent or independent variables in the top and bottom two percent observations are removed to reduce the effects of outliers. The variables are as defined in the notes to Tables 1 and 2. Summary means across the annual regressions and the related Fama and MacBeth (1973) t-statistics are provided. The last column of Panel A contains the difference between estimates of expected return from the estimation of regression (4) using I/B/E/S consensus forecasts and regression (5) using current accounting data. The last two columns of Panel B contain the differences between perfect foresight estimates and the estimates of expected return from the estimation of regression (4) using I/B/E/S consensus forecasts and regression (5) using current accounting data. Panel C repeats the analysis performed in Panels A and B using an alternative definition of price. Instead of measuring price at trade close the day after the earnings announcement, price is measured at trade close the day following the consensus forecast. This results in a price variable measured 14 days to a month and a half later. All other variables remain unchanged.

Table 4: Value-weighting observations, results of comparison of implied expected rates of return based on I/B/E/S forecasts of earnings, based on current accounting data and based on future realizations of earnings

Panel A: Descriptive statistics

Mean of annual means	Decile of market capitalization at time t									
	1 st	2 nd	3 rd	4 th	5 th	6 th	7 th	8 th	9 th	10 th
$ FE_{jt+1} $	0.419	0.397	0.398	0.443	0.428	0.455	0.466	0.488	0.579	2.369
$ FE_{jt+1} /p_{jt}$	0.102	0.053	0.040	0.034	0.026	0.023	0.018	0.017	0.015	0.012
FE_{jt+1}	-0.284	-0.235	-0.242	-0.266	-0.233	-0.237	-0.214	-0.246	-0.273	-0.890
FE_{jt+1}/p_{jt}	-0.075	-0.033	-0.025	-0.021	-0.015	-0.013	-0.009	-0.009	-0.007	-0.005
$eps_{jt+1}^{Cons}/bps_{jt}$	0.065	0.081	0.093	0.095	0.113	0.128	0.140	0.149	0.160	0.186
eps_{jt}/bps_{jt-1}	0.002	0.050	0.066	0.075	0.095	0.113	0.126	0.134	0.145	0.168
p'_{jt}/bps_{jt}	1.707	1.954	2.188	2.362	2.482	2.676	2.794	2.895	2.941	3.593
$(p'_{jt} - bps_{jt}^*)/bps_{jt-1}$	0.641	1.000	1.275	1.533	1.752	1.958	2.083	2.142	2.146	2.732

Mean of annual medians	Decile of market capitalization at time t									
	1 st	2 nd	3 rd	4 th	5 th	6 th	7 th	8 th	9 th	10 th
$ FE_{jt+1} $	0.218	0.200	0.211	0.225	0.225	0.221	0.238	0.223	0.242	0.246
$ FE_{jt+1} /p_{jt}$	0.042	0.024	0.018	0.016	0.012	0.010	0.009	0.008	0.007	0.006
FE_{jt+1}	-0.116	-0.106	-0.108	-0.116	-0.098	-0.092	-0.092	-0.090	-0.075	-0.086
FE_{jt+1}/p_{jt}	-0.023	-0.012	-0.009	-0.007	-0.005	-0.004	-0.004	-0.003	-0.002	-0.002
$eps_{jt+1}^{Cons}/bps_{jt}$	0.095	0.110	0.115	0.118	0.126	0.134	0.143	0.148	0.155	0.176
eps_{jt}/bps_{jt-1}	0.052	0.086	0.097	0.104	0.114	0.125	0.131	0.136	0.142	0.160
p'_{jt}/bps_{jt}	1.316	1.577	1.748	1.836	1.926	2.060	2.183	2.221	2.304	2.829
$(p'_{jt} - bps_{jt}^*)/bps_{jt-1}$	0.259	0.605	0.818	0.944	1.017	1.220	1.327	1.313	1.439	1.934

Table 4: Continued

Panel B: Value-weighted estimates of expected rate of return based on analysts' forecasts and current accounting data

$$\frac{eps_{jt+1}^{Cons}}{bps_{jt}} = \gamma_0 + \gamma_1 \frac{p'_{jt}}{bps_{jt}} + \mu_{jt} \quad (4)$$

$$\frac{eps_{jt}}{bps_{jt-1}} = \delta_0 + \delta_1 \frac{p'_{jt} - bps_{jt}^*}{bps_{jt-1}} + \zeta_{jt} \quad (5)$$

Analysts' consensus earnings forecasts						Current accounting data				Difference in expected rate of return
<i>T</i>	N	γ_0	γ_1	Adj R ²	$\hat{r} = \gamma_0 + \gamma_1$	δ_0	δ_1	Adj R ²	$\hat{r} = \delta_0$	
12/92	1,418	0.047 (14.73)	0.047 (44.03)	57.76%	9.35%	0.062 (23.49)	0.044 (35.38)	46.89%	6.22%	3.13%
12/93	1,544	0.052 (14.70)	0.047 (40.70)	51.76%	9.82%	0.079 (29.00)	0.042 (36.43)	46.23%	7.87%	1.95%
12/94	1,781	0.072 (22.46)	0.049 (43.95)	52.03%	12.15%	0.084 (34.82)	0.050 (48.64)	57.05%	8.39%	3.76%
12/95	1,938	0.092 (26.96)	0.036 (41.36)	46.89%	12.76%	0.127 (41.25)	0.028 (30.46)	32.37%	12.65%	0.11%
12/96	2,006	0.081 (25.50)	0.034 (45.77)	51.09%	11.53%	0.106 (38.36)	0.029 (40.29)	44.72%	10.64%	0.89%
12/97	2,137	0.094 (28.17)	0.026 (41.48)	44.60%	12.01%	0.106 (41.10)	0.023 (37.67)	39.89%	10.58%	1.43%
12/98	2,044	0.093 (28.30)	0.022 (42.72)	47.17%	11.49%	0.090 (33.70)	0.022 (45.20)	49.99%	8.97%	2.52%
12/99	1,855	0.147 (35.74)	0.010 (23.92)	23.55%	15.69%	0.147 (36.07)	0.004 (8.85)	4.00%	14.66%	1.03%
12/00	1,729	0.091 (22.09)	0.022 (36.13)	43.02%	11.26%	0.110 (28.77)	0.021 (29.60)	33.61%	11.04%	0.22%
12/01	1,808	0.059 (15.74)	0.031 (38.34)	44.84%	8.98%	0.070 (22.45)	0.030 (40.29)	47.31%	6.98%	2.00%
12/02	1,825	0.055 (18.77)	0.043 (52.26)	59.95%	9.76%	0.083 (34.75)	0.041 (54.05)	61.56%	8.26%	1.50%
12/03	2,000	0.072 (21.58)	0.032 (39.02)	43.22%	10.41%	0.098 (27.36)	0.031 (36.65)	40.17%	9.76%	0.65%
Means	1,841	0.079	0.033	47.16%	11.27%	0.097	0.030	41.98%	9.67%	1.60%
t-Statistics		(10.09)	(9.62)		(21.20)	(13.90)	(8.38)		(13.90)	(4.91)

Notes to Table 4:

Panel A of the table reports the summary statistics from repeating the analysis performed in Tables 1 and 2 by annual decile of market capitalization at time t . Panel B repeats the analysis in Table 3 using weighted least squares regression with regression weights equal to market capitalization at time t .

Table 5: Variation in the implied expected rate of return with changes in the percentage of analysts' making "buy" recommendation – minimum of five analysts following firm

Panel A: Descriptive statistics by percent of buy recommendations

	<u>90 ≤ % Buy ≤ 100</u>		<u>70 ≤ % Buy ≤ 90</u>		<u>50 ≤ % Buy < 70</u>		<u>30 ≤ % Buy < 50</u>		<u>0 ≤ % Buy < 30</u>	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median
$ FE_{jt+1} $	0.437	0.218	0.932	0.232	0.497	0.220	0.540	0.235	0.536	0.229
$ FE_{jt+1} /p_{jt}$	0.017	0.008	0.017	0.008	0.019	0.008	0.026	0.010	0.041	0.011
FE_{jt+1}	-0.268	-0.101	-0.725	-0.103	-0.251	-0.083	-0.271	-0.089	-0.287	-0.082
FE_{jt+1}/p_{jt}	-0.010	-0.004	-0.009	-0.003	-0.010	-0.003	-0.016	-0.004	-0.027	-0.004
$eps_{jt+1}^{Cons}/bps_{jt}$	0.140	0.157	0.164	0.162	0.159	0.153	0.134	0.131	0.108	0.112
eps_{jt}/bps_{jt-1}	0.125	0.150	0.152	0.151	0.143	0.140	0.120	0.120	0.091	0.101
p'_{jt}/bps_{jt}	3.860	3.011	3.435	2.686	2.848	2.305	2.371	1.921	2.029	1.649
$(p'_{jt} - bps_{jt}^*)/bps_{jt-1}$	3.649	2.313	2.844	1.948	2.005	1.438	1.485	1.016	1.032	0.704
# of observations	135		227		263		176		154	

Table 5: Continued

Panel B: Summary of results of estimation by percent of buy recommendations

$$\frac{eps_{jt+1}^{Cons}}{bps_{jt}} = \gamma_0 + \gamma_1 \frac{p'_{jt}}{bps_{jt}} + \mu_{jt} \quad (4)$$

$$\frac{eps_{jt}}{bps_{jt-1}} = \delta_0 + \delta_1 \frac{p'_{jt} - bps_{jt}^*}{bps_{jt-1}} + \zeta_{jt} \quad (5)$$

Recommendation	N	Analysts' consensus earnings forecasts				Current accounting data			
		γ_0	γ_1	Adj R ²	$\hat{r} = \gamma_0 + \gamma_1$	δ_0	δ_1	Adj R ²	$\hat{r} = \delta_0$
90 ≤ % Buy ≤ 100	135	0.100 (7.93)	0.012 (3.32)	7.90%	11.20% (9.93)	0.109 (5.12)	0.012 (1.46)	18.18%	10.94% (5.12)
70 ≤ % Buy ≤ 90	227	0.098 (9.87)	0.021 (7.73)	16.82%	11.84% (14.29)	0.102 (10.23)	0.020 (5.88)	17.42%	10.22% (10.23)
50 ≤ % Buy < 70	263	0.080 (13.67)	0.029 (12.69)	34.28%	10.82% (20.84)	0.089 (18.09)	0.028 (10.96)	30.29%	8.90% (18.09)
30 ≤ % Buy < 50	176	0.060 (7.04)	0.031 (6.80)	28.31%	9.18% (16.25)	0.072 (13.25)	0.033 (8.38)	26.85%	7.23% (13.25)
0 ≤ % Buy < 30	154	0.032 (3.13)	0.037 (9.60)	32.00%	6.86% (8.85)	0.046 (5.60)	0.044 (9.67)	30.09%	4.60% (5.60)

Table 5: Continued

Panel C: Mean differences in (t-statistics for) estimates of expected rate of return

		Analysts' expected rate of return					Expected rate of return based on current accounting data			
		$90 \leq \% \leq 100$	$70 \leq \% \leq 90$	$50 \leq \% < 70$	$30 \leq \% < 50$	$0 \leq \% < 30$	$90 \leq \% \leq 100$	$70 \leq \% \leq 90$	$50 \leq \% < 70$	$30 \leq \% < 50$
Analysts' expected rate of return	$70 \leq \% \leq 90$	-0.64% (-0.79)								
	$50 \leq \% < 70$	0.38% (0.50)	1.02% (2.11)							
	$30 \leq \% < 50$	2.02% (2.50)	2.66% (4.76)	1.64% (3.96)						
	$0 \leq \% < 30$	4.34% (5.46)	4.97% (9.01)	3.96% (8.90)	2.31% (5.04)					
Expected rate of return based on current accounting data	$90 \leq \% \leq 100$	0.26% (0.15)								
	$70 \leq \% \leq 90$		1.61% (3.14)				0.72% (0.30)			
	$50 \leq \% < 70$			1.92% (5.04)			2.04% (1.03)	1.32% (1.81)		
	$30 \leq \% < 50$				1.95% (6.38)		3.72% (1.82)	3.00% (4.77)	1.68% (3.96)	
	$0 \leq \% < 30$					2.27% (7.15)	6.35% (3.15)	5.63% (8.25)	4.31% (7.40)	2.63% (5.29)

Table 5: Continued

Notes to Table 5:

Using the median consensus analysts' forecast and the percent of buy recommendations from the summary I/B/E/S database, we estimate expected rate of return by percentage of buy recommendations for all firms with at least five analysts included in the consensus. Panel A reports descriptive statistics by percentage of buy recommendations. The variables are as defined in the notes to Tables 1 and 2. Panel B reports the results of estimating regression (4) using I/B/E/S consensus forecasts and regression (5) using current accounting data cross-sectionally using all available observations of that percentage of buy recommendations. Within the percentage of buy recommendations, observations with any of the dependent or independent variables in the top and bottom two percent observations are removed to reduce the effects of outliers. The reported numbers are the summary means across the annual regressions and the related Fama and Macbeth (1973) t-statistics. The last column for each regression in Panel B reports the annual estimates of expected rate of return by percentage of buy recommendations. Panel C reports summary means of the differences in estimates across the annual regressions and the related Fama and Macbeth (1973) t-statistics.

REFERENCES

- Baginski, S. and J. Wahlen. 2003. Residual income, risk, intrinsic values, and share prices. *The Accounting Review* 78: 327-351.
- Barber, B, R. Lehavy, M. McNicholls, and B. Trueman. 2001. Can investors profit from the prophets? Security analyst recommendations and stock returns. *Journal of Finance* 56: 531-563.
- Boni, L. and K. Womack. 2002. Wall Street's credibility problem: misaligned incentive and dubious fixes. *Wharton Financial Institutions Center*.
- Botosan, C. 1997. Disclosure level and the cost of equity capital. *The Accounting Review* 72: 323-349.
- Botosan, C. and M. Plumlee. 2005. Assessing alternative proxies for the expected risk premium. *The Accounting Review* 80: 21-54.
- Botosan, C., M. Plumlee, Y. Xie. 2004. The role of information precision in determining the cost of equity capital. *Review of Accounting Studies* 9: 233-259.
- Bradshaw, M. 2004. How do analysts use their earnings forecasts in generating stock recommendations? *The Accounting Review* 79: 25-50.
- Brown, L. 1993. Earnings forecasting research: Its implications for capital markets research. *International Journal of Forecasting* 9: 295-320.
- Claus, J. and J. Thomas. 2001. Equity risk premium as low as three percent? Evidence from analysts' earnings forecasts for domestic and international stocks. *Journal of Finance* 56: 1629-1666.
- Das, S., B. Levine, and K. Sivaramankrishnan. 1998. Earnings predictability and bias in analysts' earnings forecasts. *The Accounting Review* 73: 277-294.
- Daske, H. 2006. Economic benefits of adopting IFRS or US-GAAP – Have the expected costs of equity capital really decreased? *Journal of Business, Finance, and Accounting*. Forthcoming.
- Dhaliwal, D., L. Krull, O. Li, and W. Moser. 2005. Dividend taxes and implied cost of capital. *Journal of Accounting Research*. 43: 675-715.
- Dugar, A., and S. Nathan. 1995. The effect of banking relationships on financial analysts' earnings investment recommendations. *Contemporary Accounting Research* 12: 131-160.

- Eames, M., S. Glover, and J. Kennedy. 2002. The association between broker recommendations and broker-analysts' earnings forecasts. *Journal of Accounting Research* 40: 85-104.
- Easton, P. 2004. PE ratios, PEG ratios, and estimating the implied expected rate of return on equity capital. *The Accounting Review* 79: 73-96.
- Easton, P. 2006. Use of forecasts of earnings to estimate and compare cost of capital across regimes. *Journal of Business, Finance, and Accounting* 33: 374-394.
- Easton, P., and S. Monahan. 2005. An evaluation of accounting-based measures of expected returns. *The Accounting Review* 80: 501-538.
- Easton, P., G. Sommers, and M. Zmijewski. 2006. Effects of fraud on expectations of returns on equity investment and growth. University of Notre Dame. Working paper.
- Easton, P., G. Taylor, P. Shroff, and T. Sougiannis. 2002. Using forecasts of earnings to simultaneously estimate growth and the rate of return on equity investment. *Journal of Accounting Research* 40: 657-676.
- Fama, E. and K. French. 1992. The cross-section of expected returns. *Journal of Finance* 47: 427-465.
- Fama, E. and J. MacBeth. 1973. Risk, return and equilibrium – Empirical tests. *The Journal of Political Economy* 81: 607-636.
- Francis, J., I. Khurana, and R. Periera. 2005. Disclosure incentives and effects on cost of capital around the world. *Accounting Review* 80: 1125-1163.
- Francis, J., R. LaFond, P. Olsson and K. Schipper. 2004. Costs of equity and earnings attributes. *Accounting Review* 79: 967-1011.
- Frankel, R. and C. Lee. 1998. Accounting valuation, market expectation, and cross-sectional stock returns. *Journal of Accounting and Economics* 25: 283-319.
- Gebhardt, W., C. Lee, and B. Swaminathan. 2001. Towards an ex-ante cost of capital. *Journal of Accounting Research* 39: 135-176.
- Gode, D. and P. Mohanram. 2003. Inferring the cost of capital using the Ohlson-Juettner model. *Review of Accounting Studies* 8: 399-431.
- Groysberg, B., P. Healy, C. Chapman, and Y. Gui. 2006. Do buy-side analysts outperform the sell-side. *Journal of Accounting and Economics*. Forthcoming.

- Guay, W., S. Kothari, S. Shu. 2005. Properties of implied cost of capital using analysts' forecasts. University of Pennsylvania, Wharton School. Working paper.
- Hail, L. and C. Leuz. 2006. International differences in the cost of equity capital: Do legal institutions and securities regulation matter? *Journal of Accounting Research*. Forthcoming.
- Hribar, P. and N. Jenkins. 2004. The effect of accounting restatements on earnings revisions and the estimated cost of capital. *Review of Accounting Studies* 9: 337-356.
- Kutsoati, E. and D. Bernhardt. 1999. Can relative performance evaluation explain analysts' forecasts of earnings. Tufts University. Working paper.
- Lee, C., J. Myers, and B. Swaminathan. 1999. What is the intrinsic value of the Dow? *Journal of Finance* Volume 54 Issue 5: 1693-1741.
- Lin, H. 1994. Security analysts' investment recommendations. Stanford University. Ph.D. dissertation.
- Mendenhall, R. 1991. Evidence on the possible underweighting of earnings information. *Journal of Accounting Research* 29: 170-179.
- Michaely, R and K. Womack. 1999. Conflict of interest and the credibility of underwriter analyst recommendations. *Review of Financial Studies* 12: 653-686.
- O'Brien, P. 1988. Analysts' forecasts as earnings recommendations. *Journal of Accounting and Economics* 10: 53-83.
- O'Brien, P. 1993. Earnings forecasts and the macroeconomy. University of Michigan. Working Paper.
- O'Hanlon, J. and A. Steele. 2000. Estimating the equity risk premium using accounting fundamentals. *Journal of Business Finance and Accounting* 27: 1051-1084.
- Ohlson, J. 1989. Accounting earnings, book value and dividends: the theory of the clean surplus equation (part I). Reproduced in R. Brief and K. Peasnell (eds.) *Clean Surplus: A Link Between Accounting and Finance* (Garland), 167-227.
- Richardson, S., S. Teoh, and P. Wysocki. 2004. The walk-down to beatable analyst forecasts: the role of equity issuances and insider trading incentives. *Contemporary Accounting Research*, 21, 885-924.
- Williams, M. 2004. Discussion of "The role of information precision in determining the cost of equity capital." *Review of Accounting Studies* 9: 261-264.

Zitzewitz, E. 2002. Regulation fair disclosure and the private information of analysts.
Stanford University. Working paper.

Discounting the bull

Sell-side share analysis is wrong

But in reassuringly predictable ways

Dec 3rd 2016

- “SELL-SIDE” analysts, whose firms make money from trading and investment banking, are notoriously bullish. As one joke goes, stock analysts rated Enron as a “can’t miss” until it got into trouble, at which point it was lowered to a “sure thing”. Only when the company filed for bankruptcy did a few bold analysts dare to downgrade it to a “hot buy”.
- Economic research shows that there is some truth to the ribbing. The latest figures from FactSet, a financial-data provider, show that 49% of firms in the S&P 500 index of leading companies are currently rated as “buy”, 45% are rated as “hold”, and just 6% are rated as “sell”. In the past year, 30% of S&P 500 companies yielded negative returns.

Profits forecasts made more than a few months ahead have a dismal record of inaccuracy. According to Morgan Stanley, a bank, forecasts for American firms’ total annual earnings per share made in the first half of the year had to be revised down in 34 of the past 40 years. Studying their forecasts over time reveals a predictable pattern (see chart 1).

In theory, a diligent share analyst should do his own analysis—that is, by projecting a firm’s future revenue and expenses, and discounting them to the present. Such models, however, are extremely sensitive to different assumptions of growth rates. Since no one can know the future, analysts cheat.

Three statistical sins are common. Analysts can look at comparable companies to glean reasonable profits estimates, and then work backwards from their conclusions. Or they can simply echo what their peers are saying, and follow the herd. Or, most important, they can simply ask the companies they are following what their actual earnings numbers are.

Surveys conducted by Lawrence Brown of Temple University found that two-thirds of sell-side analysts found private calls with company managements to be “very useful” in making their estimates. Analysts’ need to maintain relationships with the companies they cover must colour

their projections. They are judged primarily on the accuracy of their short-term forecasts, so there is little risk in issuing flattering, if unrealistic, long-term projections. In the short run, however, they have an incentive to issue ever-so-slightly pessimistic forecasts, so companies can “beat” expectations. Since the financial crisis, company profits have exceeded short-term analyst forecasts around 70% of the time.

So are forecasts useless? Simply taking the market’s earnings figures from the previous year and multiplying by 1.07 (corresponding with the stockmarket’s long-run growth rate) can be expected to yield a more accurate forecast of profits more than a year in the future.

Yet the very predictability of the errors in analysts’ forecasts suggests they could be informative, if they are properly interpreted. Taking forecasts of S&P 500 earnings from 1985-2015, The Economist has built a simple statistical model to try to take out the bias that taints Wall Street’s prognostications. After controlling for the forecasts’ lead time and whether or not they were made during a recession, we find that even our relatively crude model can improve upon the Wall Street consensus for forecasts made more than a quarter in advance (see chart 2).

Adjusting for bias in short-term forecasts is harder. It is tempting simply to accept the errors—after all, they tend to be off by just a little. Data from Bloomberg show that the 320 S&P 500 companies that beat earnings expectations in 2015 did so only by a median of 1.4%. An alternative is to look at crowdsourcing websites such as Estimize. There punters—some amateur, and some professional—are shown Wall Street consensus estimates and asked to make their own forecasts. Estimize users beat Wall Street estimates two-thirds of time.

To some extent, judging Wall Street by its ability to make accurate predictions is silly. Harrison Hong, an economist at Columbia University, reckons that stock analysts should be viewed “more like media”. The latest forecasts aggregated by Thomson Reuters suggest that the S&P 500 will yield earnings per share of \$130.83 in 2017 and \$146.33 in 2018. According to our model, that would imply that they believe the actual numbers will be closer to \$127.85 and \$134.30. Share analysts want to tell the truth. They just like making it difficult.

This article appeared in the Finance and economics section of the print edition under the headline Discounting the bull

THE WALL STREET JOURNAL.

Study Suggests Bias in Analysts' Rosy Forecasts

By **ANDREW EDWARDS**

March 21, 2008; Page C6

Despite an economy teetering on the brink of a recession -- if not already in one -- analysts are still painting a rosy picture of earnings growth, according to a study done by Penn State's Smeal College of Business.

The report questions analysts' impartiality five years after then-New York Attorney General Eliot Spitzer forced analysts to pay \$1.5 billion in damages after finding evidence of bias.

"Wall Street analysts basically do two things: recommend stocks to buy and forecast earnings," said J. Randall Woolridge, professor of finance. "Previous studies suggest their stock recommendations do not perform well, and now we show that their long-term earnings-per-share growth-rate forecasts are excessive and upwardly biased."

The report, which examined analysts' long-term (three to five years) and one-year per-share earnings expectations from 1984 through 2006 found that companies' long-term earnings growth surpassed analysts' expectations in only two instances, and those came right after recessions.

Over the entire time period, analysts' long-term forecast earnings-per-share growth averaged 14.7%, compared with actual growth of 9.1%. One-year per-share earnings expectations were slightly more accurate: The average forecast was for 13.8% growth and the average actual growth rate was 9.8%.

"A significant factor in the upward bias in long-term earnings-rate forecasts is the reluctance of analysts to forecast" profit declines, Mr. Woolridge said. The study found that nearly one-third of all companies experienced profit drops over successive three-to-five-year periods, but analysts projected drops less than 1% of the time.

The study's authors said, "Analysts are rewarded for biased forecasts by their employers, who want them to hype stocks so that the brokerage house can garner trading commissions and win underwriting deals."

They also concluded that analysts are under pressure to hype stocks to generate trading commissions, and they often don't follow stocks they don't like.

Write to Andrew Edwards at andrew.edwards@dowjones.com

EQUITY RISK PREMIUM FORUM

NOVEMBER 8, 2001

Forum Participants

Martin L. Leibowitz (*Forum Chair*)

Robert D. Arnott
Clifford S. Asness
Ravi Bansal
John Y. Campbell
Peng Chen, CFA
Bradford Cornell
William N. Goetzmann
Brett Hammond
Campbell R. Harvey
Roger G. Ibbotson

Rajnish Mehra
Thomas K. Philips
William Reichenstein, CFA
Stephen Ross
Robert J. Shiller
Jeremy Siegel
Kevin Terhaar, CFA
Richard H. Thaler
J. Peter Williamson

Editorial Staff

Kathryn Dixon Jost, CFA
Editor

Bette Collins
Book Editor

Christine E. Kemper
Assistant Editor

Jaynee M. Dudley
Production Manager

Kathryn L. Dagostino
Production Coordinator

Kelly T. Bruton
Lois A. Carrier
Composition

Biographies

ROBERT D. ARNOTT is managing partner and CEO at First Quadrant, LP. Previously, he served as equity strategist at Salomon Brothers and as president and chief investment officer at TSA Capital Management. Mr. Arnott co-edited the first and second editions of *Active Asset Allocation* and *Style Management* and is the author of numerous articles. He has received four Graham and Dodd Awards of Excellence for articles published in the *Financial Analysts Journal* and is a member of the editorial boards of the *Journal of Portfolio Management*, *Journal of Investing*, and *Journal of Wealth Management*. Mr. Arnott holds degrees in economics, applied mathematics, and computer science from the University of California.

CLIFFORD S. ASNESS is president and managing principal at AQR Capital Management, LLC. Previously, he was a managing director and director of quantitative research at Goldman, Sachs & Company. Mr. Asness's articles on a variety of financial topics have appeared in many journals, including the *Journal of Portfolio Management*, *Financial Analysts Journal*, and *Journal of Banking & Finance*. In 2000, he received a Graham and Dodd Award of Excellence from the *Financial Analysts Journal*. Mr. Asness serves on the editorial board of the *Journal of Portfolio Management* and on the governing board of the Courant Institute of Mathematical Finance. He holds a B.S. in economics from The Wharton School at the University of Pennsylvania, a B.S. in engineering from the Moore School of Electrical Engineering, and an M.B.A. and a Ph.D. from the University of Chicago.

RAVI BANSAL is associate professor of finance at Fuqua School of Business at Duke University. Professor Bansal has published numerous articles in such finance journals as the *Journal of Finance*, *Review of Financial Studies*, and *Journal of Econometrics*. He is associate editor of the *Journal of Financial Econometrics* and *Journal of Securities Markets*. He holds an M.S. and a Ph.D. from the Carnegie Mellon University Graduate School of Industrial Administration and a B.A. and an M.A. in economics from the Delhi School of Economics at the University of Delhi, India.

JOHN Y. CAMPBELL is Otto Eckstein Professor of Applied Economics at Harvard University. Previously, he was Class of 1926 Professor of Economics and Public Affairs at Princeton University. He is the

author of numerous articles in financial journals and was awarded the Smith Breeden Prize from the *Journal of Finance*, the Roger F. Murray Prize from the Institute for Quantitative Research in Finance and the Paul A. Samuelson Certificate of Excellence from TIAA-CREF. He holds a B.A. from Corpus Christi College at the University of Oxford and a Ph.D. and an M.Phil. from Yale University.

PENG CHEN, CFA, is vice president and director of research at Ibbotson Associates. He conducts research projects on asset allocation, portfolio risk measurement, nontraditional assets, and global financial markets. His writings have appeared in various journals, including the *Journal of Portfolio Management*, *Journal of Association of American Individual Investors*, and *Consumer Interest Annual*. Mr. Chen holds a bachelor's degree in industrial management engineering from Harbin Institute of Technology and master's and doctorate degrees in consumer economics from Ohio State University.

BRADFORD CORNELL is professor of financial economics at the Anderson Graduate School of Management at the University of California at Los Angeles and president of FinEcon. Previously, Professor Cornell was vice president and director of the Securities Litigation Group at Economic Analysis Corporation. He has also taught at the California Institute of Technology, the University of Southern California, and the University of Arizona. Professor Cornell has served as an associate editor for a variety of scholarly and business journals and is the author of several books and numerous articles published in various finance journals. Professor Cornell holds an A.B. in physics, philosophy, and psychology, an M.S. in statistics, and a Ph.D. in financial economics from Stanford University.

WILLIAM N. GOETZMANN is Edwin J. Beinecke Professor of Finance and Management Studies and director of the International Center for Finance at the Yale School of Management. Previously, he was associate professor of finance at Yale School of Management, assistant professor of finance at Columbia Business School, and a lecturer at Ibbotson Associates. Professor Goetzmann has also served as director of the Museum of Western Art and as a writer and producer of PBS documentaries. He is the author or co-author of numerous articles on finance and real estate and is co-editor or editorial board member for

several journals. Professor Goetzmann holds a B.A., an M.A., and an M.Phil. from Yale College, an M.P.P.M. from Yale School of Management, and a Ph.D. in operations research with a specialty in finance from Yale University.

BRETT HAMMOND is director of corporate projects at TIAA-CREF Investment Management, LLC.

CAMPBELL R. HARVEY is J. Paul Sticht Professor of International Business at the Fuqua School of Business at Duke University and a research associate for the National Bureau of Economic Research. Professor Harvey is an internationally recognized expert in portfolio management and global risk management. His work on the implications of changing risk and the dynamics of risk premiums for tactical asset allocation has been published in the top academic and practitioner journals. He is editor of the *Review of Financial Studies*, co-editor of the *Emerging Markets Review*, and an associate editor for a number of journals. Professor Harvey received the 1993–1994 Batterymarch Fellowship and has received four Graham and Dodd Awards for Excellence in financial writing from the *Financial Analysts Journal*. He holds a B.A. from Trinity College at the University of Toronto, an M.B.A. from York University, and a Ph.D. from the University of Chicago.

ROGER G. IBBOTSON is chair and founder of Ibbotson Associates and the author of numerous books and articles, including the annual *Stocks, Bonds, Bills, and Inflation Yearbook*. Mr. Ibbotson is a frequent lecturer at universities, academic conferences, business conferences, and other forums. He is associate editor of the *Journal of Applied Corporate Finance* and a member of the editorial board of the *Financial Analysts Journal*. He received a *Financial Analysts Journal* Graham and Dodd Award of Excellence in 1980, 1982, 1984 and 2000. Mr. Ibbotson holds a bachelor's degree in mathematics from Purdue University, a master's degree in finance from Indiana University, and a Ph.D. in finance and economics from the University of Chicago Graduate School of Business, where he taught for more than 10 years and served as executive director of the Center for Research in Security Prices.

MARTIN L. LEIBOWITZ is vice chair and chief investment officer at TIAA-CREF Investment Management, LLC, where he is responsible for the overall management of all TIAA-CREF investments. Previously, he was a managing director, director of research, and member of the executive committee with Salomon Brothers. Mr. Leibowitz has authored several books and more than 130 articles, 9 of which

have received a *Financial Analysts Journal* Graham and Dodd Award of Excellence. He was the recipient of AIMR's Nicholas Molodovsky Award in 1995 and the James R. Vertin Award in 1998. In 1995, he received the Distinguished Public Service Award from the Public Securities Association and became the first inductee into the Fixed Income Analysts Society's Hall of Fame. Mr. Leibowitz is a trustee of the Carnegie Corporation, the Institute for Advanced Study at Princeton, and the Research Foundation of AIMR. He holds a B.A. and an M.S. from the University of Chicago and a Ph.D. in mathematics from the Courant Institute of New York University.

RAJNISH MEHRA is professor of finance at the University of California at Santa Barbara and visiting professor of finance at the University of Chicago Graduate School of Business. Previously, he was visiting professor of finance at the Sloan School of Management at the Massachusetts Institute of Technology and assistant and associate professor at the Graduate School of Business at Columbia University. He is the author of numerous articles published in such journals as *Econometrica*, the *Review of Economic Studies*, and the *Journal of Monetary Economics*, and his writings have been included in several books. He holds a B.Tech. in electrical engineering from the Indian Institute of Technology in Kanpur, India, an M.S. in computer science from Rice University, and an M.S. in industrial administration and a Ph.D. in finance from the Graduate School of Industrial Administration at Carnegie-Mellon University.

THOMAS K. PHILIPS is chief investment officer at Paradigm Asset Management Company, where he is responsible for all aspects of the investment process, including the development of new products and the enhancement of existing ones. Previously, he was managing director at RogersCasey and worked at IBM Corporation in research and active equity management for the IBM Thomas J. Watson Research Center and the IBM Retirement Fund. Mr. Philips is the author of several articles and book chapters on topics in finance, engineering, and mathematics. He holds an M.S. and a Ph.D. in electrical and computer engineering from the University of Massachusetts at Amherst.

WILLIAM REICHENSTEIN, CFA, holds the Pat and Thomas R. Powers Chair in Investment Management at Baylor University and is a consultant to TIAA-CREF. He is an associate editor of the *Journal of Investing*, on the editorial board of *Journal of Financial Education*, on the editorial review board of the *Journal of Financial Planning*, and on the editorial advisory board of the *Journal of Applied Business*

Research. Professor Reichenstein has written more than 50 articles for professional and academic journals and is a frequent contributor to the *Financial Analysts Journal*, *Journal of Portfolio Management*, and *Journal of Investing*. His work has been discussed in the *Wall Street Journal*, *Barron's*, *Forbes*, *Smart Money*, and elsewhere, and he is frequently quoted in newspapers throughout the United States. Professor Reichenstein holds a B.A. in mathematics from St. Edward's University and a Ph.D. in economics from the University of Notre Dame.

STEPHEN ROSS is Franco Modigliani Professor of Financial Economics at the Sloan School of Management at the Massachusetts Institute of Technology. He is also a principal of Roll and Ross Asset Management Corporation. Professor Ross, a widely published author in finance and economics, is best known as the inventor of Arbitrage Pricing Theory and the Economic Theory of Agency and as the co-discoverer of risk-neutral pricing and the binomial model for pricing derivatives. He is the author of *Corporate Finance*, which is in its fourth edition.

ROBERT J. SHILLER is Stanley B. Resor Professor of Economics at the Cowles Foundation for Research in Economics at Yale University. He is co-founder of Case Shiller Weiss, an economics research and information firm, and also a co-founder of Macro Securities Research, a firm that promotes securitization of unusual risks. Professor Shiller is the author of numerous articles on financial markets, behavioral economics, macroeconomics, real estate, statistical methods, and public attitudes, opinions, and moral judgments regarding markets. He is the author of *Market Volatility*, a mathematical and behavioral analysis of price fluctuations in speculative markets, and *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*, which won the 1996 Paul A. Samuelson Award from TIAA-CREF. Professor Shiller's recent book, *Irrational Exuberance*, won the Commonfund Prize for 2000. He holds a Ph.D. in economics from the Massachusetts Institute of Technology.

JEREMY SIEGEL is Russell E. Palmer Professor of Finance at The Wharton School of the University of Pennsylvania, where he has been a member of the faculty since 1976. Previously, he taught at the University of Chicago Graduate School of Business. Professor Siegel has received many awards and citations for his research and excellence in teaching, including a Graham and Dodd Award of Excellence from the *Financial Analysts Journal* and the 2000

Peter Bernstein and Frank Fabozzi Award from the *Journal of Portfolio Management*. In 1994, he received the highest teaching rating in a worldwide ranking of business school professors conducted by *Business-Week*. Professor Siegel has published articles in numerous finance journals and is the author of the book *Stocks for the Long Run* and co-author of *Revolution on Wall Street*. He holds a B.A. from Columbia University and a Ph.D. from the Massachusetts Institute of Technology.

KEVIN TERHAAR, CFA, is director of risk management and specialized investments at Brinson Partners. Previously, Mr. Terhaar was manager of investments for a trust office, where his responsibilities included asset allocation and investment analysis. He has published articles on a number of investment topics and is the co-author of "Maintaining Consistent Global Asset Views," which appeared in the *Financial Analysts Journal* in 1998. Mr. Terhaar holds a B.A. and an M.A. from the University of Virginia.

RICHARD H. THALER is Robert P. Gwinn Professor of Behavioral Science and Economics and director of the Center for Decision Research at the University of Chicago Graduate School of Business. He is also a research associate at the National Bureau of Economic Research and co-director, with Robert Shiller of the Behavioral Economics Project, funded by the Russell Sage Foundation. He has published articles on finance in such journals as the *Journal of Economic Perspectives*, *Journal of Finance*, *Journal of Portfolio Management*, and *American Economic Review*. Professor Thaler holds a B.A. from Case Western Reserve University and an M.A. and a Ph.D. from the University of Rochester.

J. PETER WILLIAMSON is the Laurence F. Whittemore Professor of Finance, Emeritus, at the Amos Tuck School of Business Administration at Dartmouth College, where he has been a member of the faculty since 1961. He is the author of more than a dozen books and monographs and numerous articles. Professor Williamson has served as a consultant and as an expert witness on cost of capital and capital structure in approximately 100 federal and state cases. He has prepared for publication summaries of all the presentations made at the semi-annual seminars of the Institute for Quantitative Research in Finance for 26 years. Professor Williamson holds a B.A. from the University of Toronto and an M.B.A., an LL.B., and a D.B.A. from Harvard University.

INTRODUCTION

Martin L. Leibowitz (*Forum Chair*)

*TIAA-CREF
New York City*

Our goal here today is to foster a very candid discussion of the many facets of the equity risk premium. Generally, the risk premium is thought of as the incremental return of certain equity market components relative to certain fixed-income components. Even when these two measures are clarified, however, which they often are not, considerable ambiguity can remain as to just what we're talking about when we talk about the risk premium. Are we talking about a premium that has been historically achieved, a premium that is the ongoing expectation of market participants, an analytically determined forecast for the market, or a threshold measure of required return to compensate for a perceived level of risk? All of these measures can be further parsed out as reflections of the broad market consensus, the opinions of a particular individual or institution, or the views of various market cohorts looking at specific and very different time horizons.

As for the issue of the risk premium as uncertainty, we often see the risk premium defined as an extrapolation of historical volatility and then treated as some sort of stable parameter over time. A more comprehensive (and more difficult) approach might be to view the risk premium as a sufficient statistic unto itself, a central value that is tightly embedded in an overall distribution of incremental returns. From this vantage point, we would then look at the entire risk premium distribution as an integrated dynamic, one that continually reshapes itself as the market evolves.

With the enormous variety of definitions and interpretations, the risk premium may seem to be the ultimate "multicultural" parameter and our forum today may have the character of a masked ball within the Tower of Babel. However, every one of us here does know and understand the particular aspect of the risk premium that we are addressing in our work. And I hope that we can communicate that clarity even as we tackle the many thorny questions that surround this subject. The risk premium is a concept that is so central to our field of endeavor that it might properly be called the financial equivalent of a cosmological concept.

Theoretical Foundations I

Richard H. Thaler

*Graduate School of Business
University of Chicago
Chicago*

One of the puzzles about the equity risk premium is that in the U.S. market, the premium has historically been much greater than standard finance theory would predict. The cause may lie in the mismatch between the actual asset allocation decisions of investors and their forecasts for the equity risk premium. In this review of the theoretical explanations for this puzzle, two questions are paramount: (1) How well does the explanatory theory explain the data? (2) Are the behavioral assumptions consistent with experimental and other evidence about actual behavior? The answers to both questions support the theory of “myopic loss aversion”—in which investors are excessively concerned about short-term losses and exhibit willingness to bear risk based on their most recent market experiences.

A good place to start consideration of what the equity risk premium should theoretically be is a discussion of the risk premium puzzle: The equity risk premium in the U.S. market has historically been much bigger than standard finance theory would predict. Based on the familiar Ibbotson Associates (2001) data of the long-term historical return to U.S. stocks, T-bonds, and T-bills, if you had invested \$1 in the stock market at the end of 1925 (with dividends reinvested), you would now have more than \$2,500; if you had put \$1 in T-bonds, you

would have about \$49; and if you had put \$1 in T-bills, you would have only \$17. These differences are much too large to be explained by any reasonable level of risk aversion.

The Puzzle

The formal puzzle, which was posed by Mehra and Prescott (1985), is that, on the one hand, if you ask, “How big a risk premium should we expect?” the standard economic model (assuming expected-utility-maximizing investors with standard additively separable preferences and constant relative risk aversion, A) provides a much smaller number than is historically true, but if you ask, “How risk averse would investors have to be to demand the equity risk premium we have seen?” (that is, how large does A have to be to explain the historical equity premium), the answer is a very large number—about 30. Mehra and Prescott’s response was that 30 is too large a number to be plausible.

Why? What does a coefficient of relative risk aversion of 30 mean? If I proposed to you a gamble in which you have a 50 percent chance that your wealth will double and a 50 percent chance that your wealth will fall by half, how much would you pay to avoid the chance that you will lose half your wealth? If you have a coefficient of relative risk aversion equal to 30, you would pay 49 percent of your wealth to avoid a chance of losing half your wealth, which is ridiculous. And that is why I believe that investors do not have a coefficient of relative risk aversion of 30.

Another way to think about this puzzle is that for reasonable parameters (and theorists argue about what those are), we would expect an annual risk premium for stocks over bonds of 0.1 percent (10 basis points).

In the Mehra–Prescott model, the coefficient of relative risk aversion, A , is also the inverse of the elasticity of intertemporal substitution, so a high value of A implies an extreme unwillingness to substitute consumption tomorrow for consumption today, which implies a long regime of high interest rates. We have not, however, observed high interest

rates for extended periods of time. Historically, the risk-free rate has been low, barely positive for much of the 20th century. Therefore, part of the risk premium puzzle is the “risk-free-rate puzzle”: Why do we not see very high interest rates if investors are so risk averse?

How do we resolve these puzzles? One answer is to “blame the data”—for example, survivorship bias. The returns in the U.S. equity market have been particularly favorable, which may be simply the product of good luck. In other words, some markets have collapsed and disappeared. So, we should not focus all our attention on one market in one period; one market can go awry.

My view is that if we can worry about stock markets going awry, we had better also worry about bond markets going awry. For example, over the long run, bond investors have experienced bad periods of hyperinflation. Bond investors have been wiped out by hyperinflation just as stock investors have been wiped out by crashes. So, if we are going to consider the effect of survivorship bias on the data, we need to look at both sides of the equation—stock and bond returns—which brings us back to a puzzle. If you adjust *both* returns for risk, you still end up with a puzzle.

The part of the puzzle that I want to stress is the contrast between investor investments and investor expectations. I am a behaviorist, and the behavior I find puzzling is how investor expectations fit with their investments.

Throughout the 1980s and 1990s, investors had expectations of a big equity premium, typically in the range of 4 percent to 7 percent. **Table 1** provides the results of a survey of fund managers on their forecasts for U.S. security returns at two points in time almost 10 years apart. Note that investor estimates of the equity risk premium fall into the 4–6 percent range in both years.

Other evidence comes from surveys of forecasts of the 10-year equity risk premium over the last decade (for example, Welch’s 2000 survey of econo-

mists); again, the estimates are substantial. A problem with such surveys, of course, is that we never know the question the people were really answering. For example, most respondents, including economists, do not know the difference between the arithmetic and the geometric return, and this confusion can skew the results. So, we cannot know precisely what such surveys show, but we can know that the estimates of the equity risk premium are big numbers compared with an estimate of 0.1 percent.

Thaler’s Equity Premium Puzzle

The real puzzle is a mismatch between the allocations of investors and their forecasts for the equity risk premium. Many long-term investors—individuals saving for retirement, endowments, and pension fund managers—think the long-term equity risk premium is 4–5 percent or higher yet still invest 40 percent of their wealth in bonds. This phenomenon is the real puzzle.

One version of this puzzle is “Leibowitz’s Lament.” In a former life, Marty Leibowitz was a bond guy at Salomon Brothers. As a bond guy, his job was to give investors a reason to buy bonds. The numbers Marty was crunching in 1989 for the wealth produced by \$1 in stocks versus the wealth produced by \$1 invested in bonds could have been those from the Ibbotson Associates studies. The historical risk premium was 6.8 percent, which made the return numbers ridiculous. Marty’s analysis showed that if we assume investors may lever, the correct asset allocation at that time would have been at least 150 percent in equities. The puzzle is that investors did not invest this way then and do not do so now.

Theoretical Explanations

Many explanations for the puzzle have been offered, and all the theoretical explanations so far proposed are behavioral—in the sense that they build on the Mehra–Prescott model and then make some inference about investor preferences. In most of these models, the investors make rational choices but their preferences are still slightly different from ones traditionally considered normal.

Epstein and Zin (1989) broke the link that A is equal to the coefficient of relative risk aversion and the elasticity of intertemporal substitution. With their approach, the standard assumptions of expected utility maximization are destroyed.

Constantinides (1990) introduced the theory of habit formation based on the following postulate: If I’m rich today, then I’m more miserable being poor tomorrow than if I’d always been poor. A similar theory of habit formation, the approach of Abel (1990), is based on the concept of “keeping up with

Table 1. Forecasted Returns: Survey of Fund Managers
($N = 395$)

Fund/Premium	1989	1997
90-day T-bills	7.4 %	4.7 %
Bonds	9.2	6.9
S&P 500	11.5	10.4
S&P 500 – T-bills	4.1	5.7

Source: Greenwich Associates.

the Joneses.” Perhaps the leading model at the moment, however, is that of Campbell and Cochrane (1995, 1999), which combines the idea of habit formation with high levels of risk aversion. Together, these behavioral theories appear to explain some, but not all, of the data—including the risk-free-rate puzzle.

Benartzi and I (1995) suggested the theory of loss aversion, which is the idea that investors are more sensitive to market changes that are negative than to those that are positive, and the idea of mental accounting, which adds that investors are more sensitive when they are given frequent market evaluations. Combined, loss aversion and mental accounting produce what we called “myopic loss aversion.” We explicitly modeled investors as being myopic, in that they think about and care most strongly about the market changes that occur over short periods, such as a year.

Barberis, Huang, and Santos (1996) used the myopic loss aversion model and added another behavioral phenomenon, the “house money effect” (that is, loss aversion is reduced following recent gains), in an equilibrium model. When people are ahead in whatever game they are playing, they seem to be more willing to take risks. I also documented this effect in some experimental work about 10 years ago. I discovered this phenomenon playing poker. If you’re playing with people who have won a lot of money earlier in the game, there is no point in trying to bluff them. They are in that hand to stay.

So, we have a long list of possible behavioral explanations for the equity risk premium. How do we choose from them? We should concentrate on two factors. The first factor is how well the models fit the data. The second factor, and it is a little unusual in economics, is evidence that investors actually behave the way the modeler claims they are behaving. On both counts, the myopic loss aversion arguments that Benartzi and I (1995) proposed do well.

First, all the consumption-based models have trouble explaining the behavior of two important groups of investors, namely, pension funds and endowments. And these two groups hold a huge amount of the equity market in the United States.

Second, I do not understand why habit formation would apply to a pension-fund manager or the manager of the Rockefeller Foundation.

Third, explanations based on high levels of risk aversion do not fit the following situation: Consider these gambles. Gamble 1: You have a 50 percent chance to win \$110 and a 50 percent chance to lose \$100. Gamble 2: You have a 50 percent chance to win \$20 million and a 50 percent chance to lose \$10,000. Most people reject Gamble 1 and accept Gamble 2.

Now, those two preferences are not consistent with expected utility theory. To be consistent with expected utility theory, if you reject the first gamble, you must also reject the second gamble. This inconsistency between behavior and utility theory is a problem for all the models except those that incorporate loss aversion and “narrow framing.” In narrow framing, people treat gambles one at a time.

In Thaler, Tversky, Kahneman, and Schwartz (1997), we reported on some experiments to determine whether investors actually behave the way our myopic loss aversion model says they do. In the first experiment, we sat participants down at a terminal and told them, “You are a portfolio manager, and you get to choose between two investments, A and B.” One choice was stocks, and the other was bonds, but they were not told that. They were simply shown each investment’s returns for the investment period just completed. At the end of every period, the pseudo portfolio managers were instructed to invest their money for the upcoming period based only on the prior-period returns for A and for B. So, they made an asset allocation decision every period. The participants were paid based on the amount of wealth their portfolio had earned at the end of the experiment.

To test the effect of how often investors receive feedback, in various runs of the experiment, we manipulated “how often” the participants were able to look at the return data. In the learning period, the participants learned about the risk and returns of the investments over time. One group of participants received feedback the equivalent of every six weeks, which led to a lot of decision making. Another group made decisions only once a year. So, the first group was working in a condition of frequent evaluation, whereas the second group was receiving exactly the same random feedback as the first one but the returns for the first eight periods were collapsed into a single return. A third group was given a five-year condition. We also had an “inflated monthly” condition in which we increased returns by a constant over the 25-year period that was sufficient to create periods with never any losses. Over the 25 years, 200 decisions were being made in the most frequent condition and 5 in the least frequent condition.

When that part of the experiment was completed and the participants had enjoyed plenty of opportunity to learn the distribution patterns, we instructed them to make one final decision for the next 40 years. Outcomes were “yoked” to assure that all manipulations had the same investment experience.

Our hypotheses were, first, that more frequent reports would induce more risk aversion, resulting in an increased allocation to bonds and, second, that shifting the returns of both assets up to eliminate

losses would make stocks relatively more attractive. **Table 2** presents the results.

Table 2. Effect of Frequency of Feedback: Allocation to Bonds		
Feedback Group	Number	Mean
<i>A. Final decisions</i>		
Monthly	21	59.1 %
Yearly	22	30.4
Five year	22	33.8
Inflated monthly	21	27.6
<i>B. Decisions during the last five “years”</i>		
Monthly	840	55.0 %
Yearly	110	30.7
Five year	22	28.6
Inflated monthly	840	39.9

As you can see, participants involved in the monthly condition (the most frequent decision-making condition), on average, chose to invest 60 percent of their money in bonds. Participants in the yearly condition chose to invest only 30 percent in bonds. The participants made the most money if they chose 100 percent stocks every period.


We concluded that the more often investors look at the market, the more risk averse they become, which is exactly what our theory suggests. Loss aversion can be mitigated by forced aggregation (to avoid narrow framing), and learning may be improved by less frequent feedback.

Another set of experiments on myopic loss aversion involved 401(k) participants—specifically, staff among University of Southern California employees who had become eligible for the program in the past year. They were shown return data for Fund A (pro-

viding higher returns than Fund B but riskier, equivalent to stocks) and Fund B (equivalent to bonds) and then asked how they would allocate their money. One group was given one-year returns and one group was given 30-year returns. **Figure 1** contains the charts presented in which the historical equity risk premium was used. The figure shows the distribution of periodic rates of return that were drawn from the full sample. That is, if this is the distribution you’re picking from, what allocations would you make? Possible outcomes are ranked from worst on the left to best on the right. When we showed the participants the distribution of 1-year rates of return for each asset category (Panel A), the average choice was to invest about 40 percent in stocks. Stocks seemed a bit risky to participants under this scenario. When we showed exactly the same data as compounded annual rates of return for a 30-year investment (Panel B), the participants chose to put 90 percent of their money in stocks. The data are the same in both charts, but the information is presented in a different way. Again, we concluded that the amount investors are willing to invest in stocks depends on how often they look at periodic performance.


Finally, we showed participants the data with a lower risk premium. As **Figure 2** shows, we divided the equity premium in half. Again, Panel A shows the revised return data for the 1-year periods, and Panel B shows the revised return data for the 30-year period. In this experiment, the participants liked stocks equally well either way they viewed the data. They chose to put about 70 percent of their money in stocks in either scenario. We call this situation a “framing equilibrium.” If the equity premium were a number such as 3 percent, investors would put about the same amount of money into the stock market whether they had a long-term perspective or not.

Figure 1. Charts Constructed with Historical Risk Premium of Equity over Five-Year T-Bonds



Notes: Fund A was constructed from the historical returns on the NYSE value-weighted index, and Fund B was constructed from the historical returns on five-year U.S. T-bonds.

Figure 2. Charts Constructed with Half the Historical Risk Premium of Equity over Five-Year T-Bonds



Notes: Fund A was constructed from the historical returns on the NYSE value-weighted index, but 3 percentage points were deducted from the historical annual rates of return on stocks. Fund B was constructed from the historical returns on five-year U.S. T-bonds.

Theoretical Foundations I

Richard H. Thaler

*Graduate School of Business
University of Chicago
Chicago*

SUMMARY

by Peter Williamson

*Amos Tuck School of Business Administration
Dartmouth College, Hanover, New Hampshire*

Richard Thaler was the first to speak to the group and the only one dealing essentially with behavioral finance aspects of the equity risk premium puzzle.

He started by discussing the now familiar Ibbotson Associates data from the *2000 Yearbook*,¹ showing the cumulative value of a dollar invested at the end of 1925 in U.S. stocks, T-bonds, and T-bills, with the stock investment (with reinvested dividends) growing to more than \$2,500 while a dollar invested in T-bonds grew to about \$49 and one invested in T-bills to only \$17 by the year 2000. The difference, he said, is much too large to be explained by any reasonable level of risk aversion. Thaler described analysis showing that a 0.1 percent (10 basis point) per year premium for stocks over bonds would be a reasonable equilibrium risk premium; the actual excess return, however, has been more than 7 percent.

In the Mehra-Prescott (1985) model, the constant relative risk aversion, which would have to be 30 to explain the actual historical excess return of stocks, is also the inverse of the elasticity of intertemporal substitution. A value of 30 is very high and implies very high interest rates. But interest rates since 1925 have not been high enough to justify that risk aversion.

What, then, is the explanation for the high historical excess return on stocks? One possibility is high risk coupled with good luck investing in the U.S. stock market. But bond markets are risky too, and if both stock and bond returns are adjusted for high risk, we are still left with an extraordinary gap in historical returns. Furthermore, most surveys in the

1980s and 1990s of “expert” opinion indicated a high expected equity premium, on the order of 4–6 percent. And current surveys give consistent results. Thaler’s observation is that many long-term investors who think that the long-term equity premium is 4–5 percent, or higher, still invest 40 percent in bonds, something that is not easily explained. A firm belief in such a premium should have led to at least a 100 percent allocation to stocks. The size of the historical excess equity return versus the size of the expected equity premium present a puzzle.

Most attempts to explain the puzzle focus on behavioral deviations from the standard assumptions of expected utility maximization. Epstein and Zin (1989) broke the link between the coefficient of relative risk aversion and the elasticity of intertemporal substitution. Constantinides (1990) incorporated “habit formation” to posit rising risk aversion with high returns. Others see further reasons for very high risk aversion; they include Benartzi and Thaler (1995) in their myopic risk aversion model.

Thaler put forward a test for choosing among explanations in the form of two questions: (1) How well does the explanatory theory explain the data? (2) Are the behavioral assumptions consistent with experimental and other evidence about actual behavior?

The answers to both questions, he said, support the myopic loss aversion theory. All the consumption-based models have trouble explaining the behavior of pension funds and endowments. A number of experiments presenting people with choices of different gambles have argued against the high-risk-aversion theory. At the same time, experiments posing a problem of allocating funds between stocks and T-bonds have supported myopic loss aversion. Participants in these experiments were asked to allocate money between stocks and bonds after receiving periodic reports on the investment performance of the two classes. It was found that providing more frequent performance feedback induces greater risk aversion and hence reduces commitment to stocks. Shifting

¹ See Ibbotson Associates (2001).

upward and equally the reported returns for both asset classes such that there were no losses for either led to greater investment in stocks.

A further experiment asking subjects to divide retirement funds between stocks and bonds on the basis of the historical excess return on stocks led to a median 40 percent investment in stocks when the subjects were shown distributions of one-year returns and to a median 90 percent investment in stocks when the distributions shown were of 30-year returns.

When the reported excess return on stocks was cut in half from its historical level and the experiment was repeated, the median allocation to stocks was about 70 percent for the annual and for the 30-year distributions. Thaler referred to this condition as

“framing equilibrium.” The expected risk premium was now such as to remove the influence of the time period of the performance results studied. The equilibrium was reached at an equity premium of about 3 percent.

His three final conclusions were as follows:

- The historical excess return on equities has been surprisingly high.
- Part of the explanation seems to be that investors are excessively concerned about short-term losses.
- Part may be that willingness to bear risk depends on recent experience, both because past gains provide a psychological cushion against future losses and because high returns can create unrealistic expectations about the future.

Theoretical Foundations II

Clifford S. Asness

*AQR Capital Management, LLC
New York City*

Historically, high P/Es have led to low returns and low P/Es have led to high returns. So, with today's market at historically high P/Es, there is a real need for rescue. This discussion examines three possible ways in which the market might be saved from decline: high and sustained real earnings growth (which is highly unlikely), low interest rates (which help only in the short term), and investor acceptance of lower future rates of return. The last possibility boils down to a choice between low long-term returns forever and very low (crash-type) returns followed by more historically normal returns. The research presented here found some support for the prescription that investors should accept a 6–7 percent nominal stock return, but evidence indicates that investors do not actually think they are facing such low returns.

My talk does not fit neatly into the category of “theoretical foundations,” which makes sense; after all, someone who runs a hedge fund is not going to have much to add to the theoretical foundations that underlie our musings about the equity risk premium, certainly not in this crowd!

My first set of data is intended to be an icebreaker. As a beginning, **Figure 1** plots the S&P 500 Index's P/E from 1881 to 2001. From those data, I created seven P/E buckets, or ranges, covering the 1927–2001

period. For each of the buckets, I calculated the median real annualized stock market return for the following decade and the worst return for any decade. **Table 1** provides the results for each range. We can argue about statistical significance, but these numbers are pretty striking. The infallibility of stocks is typically drawn from a 20-year horizon, so I have cheated by using a 10-year horizon. But the infallibility still exists when stocks are bought at low valuation ratios.

The note “Here Be Dragons” is a caution about what might happen with those P/Es of 32.6 to 45.0. It is a saying (similar to “Terra Incognita”) once used on old maps for areas not yet visited. The highest P/E, about 45, was reached in 2000. We don't know what the next 10 years will bring. We still have another eight and a half years to go, but for the one and a half years we have recently visited, the return realization is fitting the chart nicely.


The relationship between starting P/E and subsequent return is potentially exaggerated because much of the strong relationship comes from P/E reversion. What if P/Es did not change?

Figure 2 presents some input into the relationship if P/Es were constant. In the figure, trailing 20-year real S&P earnings growth is plotted for the past 110 years. For this period, annualized real earnings growth averaged 1.5–2.0 percent fairly consistently. Those people who actually still assume 10 percent nominal returns on stocks should recognize that such a return would require 5–6 percent real earnings growth over the next 10–20 years. Such growth has happened only a few times in history, and it has happened only after very depressed market conditions, which we are not really experiencing now, certainly based on the last 10 years. With a 2 percent real earnings growth forecasted, a long-term buy-and-hold investor in the S&P 500 can expect to earn 6–7 percent nominal returns.

What Can Save the Stock Market?

I envision a bad 1920s-type serial in which the villain has tied the stock market to the railroad tracks and a

Figure 1. Historical P/E of the S&P 500, 1881–2001




Note: P/E was calculated as the current price divided by the average of earnings for the past 10 years adjusted for inflation.

Table 1. Real Stock Market Return in the Next 10 Years for Historical P/E Ranges of the S&P 500, 1927–2001 Data

P/E Range (low to high)	Median Return (annualized)	Worst Return (total)
5.6 to 10.0	11.0 %	46.1 %
10.0 to 11.7	10.6	37.3
11.7 to 14.1	10.0	4.1
14.1 to 16.7	9.0	–19.9
16.7 to 19.4	5.4	–23.1
19.4 to 32.6	–0.4	–35.5
32.6 to 45.0	Here Be Dragons!	

Figure 2. S&P 500 Trailing 20-Year Real Earnings Growth, 1891–2001



Note: Earnings growth is annualized.

voice-over is pleading, “What can save stocks?” This question is going to be the organizing principle for my presentation today. I am going to concentrate on three things that could save stocks, although other answers

may be possible. One is sustained high real earnings growth—“high” meaning better than the historical average. The second, a Wall Street favorite, is the so-called Fed model, in which the U.S. Federal Reserve

lowers interest rates and supports high P/Es. The third is a simple hero—investor acceptance of lower future rates of return in the long term.

HIGH EARNINGS GROWTH. First, something we all probably know: Only if the future brings extra-special, super-high earnings growth are very high starting P/Es justified. For each level of P/E at the start of a 10-year period except very low P/Es (when returns are always on average strong), decades with stronger earnings growth also experienced stronger average stock returns, and even when P/Es were high, if earnings growth came in very high, returns were on average strong. This analysis, however, gives us an *ex post*—not a predictive—measure. If we see extraordinarily high growth in real earnings after 2001, we will probably see high real equity returns. However, the question is: What reason do we now have to be optimistic that such abnormally high earnings growth will occur?

One reason is that higher productivity and technological advancement could create high earnings growth. I think this development is unlikely. Historically, most productivity benefits accrue to workers and consumers, not necessarily to earnings:

Optimists frequently cite higher growth of real output and enhanced productivity, enabled by the technological and communications revolution, as the source of this higher growth. Yet the long-run relationship between the growth of real output and *per share* earnings growth is quite weak on both theoretical and empirical grounds. (Siegel 1999, pp. 14–15)

So, the first hurdle to believing in high earnings growth is to believe the productivity numbers, and the second is to believe earnings will benefit.

Now, let's look at the empirical data. In **Table 2**, I show the historical relationship between P/E at the beginning of a period and subsequent average 10-year real earnings growth for 1927–2001. The numbers in the 16 quadrants, or 16 buckets, are actual realized real earnings growth over rolling 10-year periods.

Each number corresponds to a range of starting P/Es and a range of starting earnings retention rates. Historically, when both the starting P/E and the retention rate are high, the real earnings growth rate is low. On May 30, 2001, the P/E of the S&P 500 was 27.3 and the retention rate was 65.3 percent, which today puts us in the bottom right bucket, so the dragons are off to the right. This position is not promising for saving stocks.

We can interpret Table 2 further. The second way stocks could experience future high earnings growth is through market efficiency. The idea is that in an efficient market, high current P/Es will lead to higher earnings growth because the market must be right. I like this approach. I wish it were the case, but I don't think the data support it well. Table 2 shows no relationship between starting P/E and future earnings growth. In fact, P/E does a lousy job of predicting earnings growth. I will go further. It does no job. In fact, the data show that higher P/Es have not led to higher real earnings growth going forward and lower P/Es have not led to lower growth. The joint hypothesis of constant expected returns and market efficiency should lead to P/Es predicting growth, but the hypothesis doesn't hold, at least in the data.

Finally, Table 2 sheds light on the third reason we might now expect high earnings growth: the idea that high cash retention (low payout ratios) leads to strong growth. Table 2 indicates, however, that the retention rate at the beginning of a period has been *inversely* related to the subsequent 10-year growth in earnings. The impact of the retention rate is incredibly, astronomically backward. Rob Arnott and I have struggled with this phenomenon. We haven't found this impact to be intuitive—it is not a forecasted result—but we do have a few *ex post* theories as to why higher retention rates might lead to lower real growth rates. I'll share three of them quickly.

The first reason relates to company managers. The general idea is that companies retain a lot of cash

Table 2. Average 10-Year Real Earnings Growth, 1927–2001 Data

Starting P/E	Retention Rate (%)				
	Negative to 37.7	37.7 to 44.4	44.4 to 50.3	50.3 to 63.9	63.9 →
5.9 to 10.4	4.1 %	2.5 %	2.2 %	–0.3 %	
10.4 to 13.8	4.3	2.5	2.4	0.6	
13.8 to 17.2	3.3	2.5	1.7	–0.4	
17.2 to 26.3	4.3	2.7	0.8	–0.6	
26.3 →					The Dragons Are Here!

to finance projects for behavioral reasons such as empire building. If the cash is for projects, managers are not doing a good job with the cash; they tend to pursue and overinvest in marginal projects, which is reflected in the future lowered growth rates of the company. If this is the explanation, the telecom boom in the late 1990s is going to be the poster child for empire building for all eternity.

Another theory, less plausible in my opinion, is that managers have information that the market doesn't have. It is generally accepted that companies are loath to cut dividends. So, the theory goes that when a company's managers pay high dividends, the market perceives that those managers must have such positive information about the company's prospects that they know they will not have to cut dividends in the future. When managers pay high dividends, they are optimistic because they have information unknown in the market. When managers do not pay high dividends, they must be nervous. So, retention of earnings may reflect a desire by managers to smooth dividends.

The third explanation is that Rob and I are doing something wrong. We have each double-checked our approach and the data repeatedly, but when you get a wacky result, for intellectual honesty, you still have to admit the possibility. That is why I mentioned the dragons, because we are off the charts and into uncharted territory.

If history repeats and higher P/Es and higher retention rates lead to lower real earnings growth and if Rob and I are not making an error, the future does not bode well for real earnings growth.

LOW INTEREST RATES. The second possible way stocks can be saved is low interest rates. **Figure 3** compares the P/E (or the "absolute" value of the S&P 500) with the earnings yield on the S&P 500, E/P, minus the 30-year U.S. T-bond yield, Y (or the "relative" value of the S&P 500); Panel A graphs these indicators for the past 20 years. As you can see, P/E has certainly fallen from its peak in 1999 but is still at the high end of the 20-year range. The equity yield minus the bond yield is one version of the Fed model. In that model, a high value is an indication of good news for the equity market, but for P/E, a high value indicates bad news for the market. Using the Fed model, the situation does not look that bad in 2001; the market is above average on earnings yield minus bond yield.

The same information, but stretching back to 1927, is presented in Panel B of **Figure 3**. The line for earnings yield minus bond yield is pretty lackluster over the period. When stocks were far cheaper in relation to bonds, stocks used to be bought for their

dividend yield; this chart uses earnings yield, but the difference is not really important. As Panel B shows, if Wall Street had a little bit longer perspective, such as looking back to 1927 rather than just 20 years, even the Fed model, or the relative value of the equity market, does not look great.


Forgetting the data, note that the Fed model has little theoretical standing. Nominal earnings growth does correlate nicely with expected inflation over time. A lot of confounding biases, such as depreciation methods, accounting choices, and different inflationary environments, affect the P/E calculation (see Siegel 1998). But by and large, the net of those biases is not clear. What does appear fairly clear, however, is that the market does not seem to understand that if you write down the expected return of a stock (dividend yield plus earnings growth), then if inflation and interest rates fall and earnings growth drops along with them, the P/E does not have to change. I think you understand the concept, but it is an idea I have to explain to most people, and I encourage you to do the same. People believe P/Es have to move with interest rates, and they are probably wrong, or at least overstating the relationship.

Figure 4 shows a plot of the S&P 500's realized 20-year volatility divided by the bond market's 20-year realized volatility against the relative yield of the stock market for 1950 to 2001.¹ I chose 20 years because I think of 20 years as a generation, so the ratio plotted from the *x*-axis reflects what a generation thinks in terms of how risky stocks are versus bonds. This ratio is a very robust indicator for each five-year period, up to 30 years. The *y*-axis is the earnings yield on the S&P 500 minus the 10-year bond yield. Whenever you look at long-term autocorrelated relationships like this, you have to carry out many, many robustness tests. This ratio survived every test we came up with.

Note that the *y*-axis is not stock yields; it is stock yields *minus* nominal bond yields. The market clearly does trade on interest rates in the short term. Not many models have a high R^2 at forecasting short-term (less than a one-year horizon) market performance. One indicator that is less pathetic than most in this regard is deviation from the fitted [linear (normal)] line in **Figure 4**. However, for longer horizons, such as forecasting the next 10-year real stock return, neither the bond yield nor the volatility measures matter. P/E alone forecasts the real stock return. So, an investor with a short horizon cares a lot about this line, but an investor with a long horizon doesn't.

¹ Figure 4 is similar to Figures 7 and 8 in Asness (2000b). In that article, Figure 7 goes back to 1871 and forward to mid-1998 and Figure 8 goes back to 1881 and forward to mid-1998.

Figure 3. S&P 500 “Absolute” and “Relative” Value



Note: S&P 500 P/E and E/P; 10-year T-bond yield.


I have marked on Figure 4 where we were on February 28, 2000, and on September 30, 2001. On February 28, 2000, short-term traders could not be saved by anything; the solid triangle is well under the line. Stocks were yielding much less than they had historically—even given unusually low volatility and unusually low interest rates relative to the historical average.

The September 2001 mark in Figure 4 indicates that stock performance doesn't look too bad over the very short term. Short-term investors tend to trade on

this relationship—that is, trade on the idea that eventually the market moves back to the line for behavioral reasons. Note that this relationship is behavioral because it is based on errors—which does not change what the equity risk premium is in the long term. Over the short term, it is the deviation of E/P from the line that counts; over the long term, it is only the actual E/P that counts.

ACCEPTANCE OF LOW RETURNS. Now for the third possible hero that might save the stock

Figure 4. Stock versus Bond Valuation, 1950–2001



Note: S&P 500 E/P; 10-year T-bond yield.

market: Are investors willing to accept low stock returns? Have they understood the idea that future returns will be low, as so many of us have discussed. A ton of “strategists” will give explanations of why high P/Es are supportable, but then they will follow the explanations with the expectation of 10–12 percent stock returns anyway. That reasoning is questionable to say the least. The first part is believable; no one can say that a 1–2 percentage point return over bonds is bad. But you cannot have your cake and eat it too. Or as I like to say when it comes to Wall Street investors, they cannot have their cake and eat yours too.

What if investors haven’t yet realized the conundrum of expectations versus reality? Surveys exist—Campbell Harvey is going to present his survey data [see the “Implications for Asset Allocation, Portfolio Management, and Future Research” session]—that indicate respondents are expecting very high equity returns. Survey data are not always the most reliable, but the data report that the high return expectations are out there. I talk to a lot of pension plans, and not many of them are using assumptions as low as 6–7 percent nominal returns or a 1 percent real equity return over bonds. And investors who plan to retire at 38 because they expect to get a 5 percent equity risk premium and 7 percent real stock returns forever are going to wake up at 62 out of money.


Are investors rationally accepting the low equity risk premium, or are a lot of people still trying to buy lottery tickets?² Many have shown that Wall Street’s growth expectations are ridiculously optimistic, but investors seem to still believe them. So, Rob and I examined a strategy based on these expectations. We formed a portfolio for a 20-year period that was long high-growth stocks and short low-growth stocks (based on Wall Street’s estimates). **Figure 5** shows the rolling 24-month beta of that long–short portfolio from December 1983 to September 2001. For a long time, the beta was mildly positive, but for the past few years, it has been massively positive. It is a dollar long, dollar short 0.5 beta. Figure 5 says that every rally for the past several years has occurred because the high-expected-growth stocks were crushing the low-expected growth stocks. And every market sell-off has been a result of the opposite occurring. Does this pattern indicate rational acceptance of the low equity risk premium or the buying of lottery tickets?

Conclusion

Broad stock market prices are still well above those of most recorded history (and of all history excluding 1999–2000 and just before the crash of 1929). Unless a miracle happens, we must prepare for very low returns as compared with history. In the end, the market offers two choices: low long-term expected

² See Statman (2002).

Figure 5. Rolling 24-Month Beta of Long–Short Portfolio, December 1983–September 2001



Note: Except for 2001, dates are as of December.

returns in perpetuity or very bad short-term returns with higher, more normal expected returns in the long run. My personal opinion: Do the events of

1999–2001 strike anyone as a group of rational investors embracing and accepting a permanently low risk premium? If so, I missed it on CNBC.

Theoretical Foundations II

Clifford S. Asness

*AQR Capital Management, LLC
New York City*

SUMMARY

by Peter Williamson

*Amos Tuck School of Business Administration
Dartmouth College, Hanover, New Hampshire*

Clifford Asness made the second presentation of the day, beginning with a graph (**Figure 1**) showing the record of the S&P 500 Index's P/E (current price divided by the average of the preceding 10 years' real earnings) for 1881 to 2001. The highest P/E, about 45, was reached in 2000. **Table 1** reports for each of six ranges of P/E the median real stock market return in the next 10 years and the return for the worst decade. In general, high P/Es led to low subsequent returns and to the worst of the worst decades and low P/Es led to high returns and to the best of the worst decades.

Asness observed that much of what Table 1 shows in terms of consequences of P/E levels comes from P/E reversion. Some would ask: What happens if the ratios do not revert? **Figure 2**, showing S&P 500

trailing 20-year real earnings growth (annualized) helps to answer the question.

Asness next examined three possible ways in which the market might be saved from decline. One is high and sustained real earnings growth. A second (the Wall Street solution) is low interest rates. This is the so-called Fed model. The third way is based on investor acceptance of lower future rates of return. This answer would mean no imminent crash but a less attractive long-term return.

Would high earnings growth work? **Table 2** shows the historical relationship between P/E at the beginning of a period and subsequent average 10-year real earnings growth for 1927–2001. The numbers in the 16 quadrants, or 16 buckets, are actual realized real earnings growth over rolling 10-year periods. Each number corresponds to a range of starting P/Es and a range of starting earnings retention rates. Historically, when both the starting P/E and the retention rate are high, the real earnings growth rate is low.

Why might we expect high earnings growth? Some might say because of increasing productivity and technological advancement. But the relationship between growth of real output and *per share earnings*

Figure 1. Historical P/E of the S&P 500, 1881–2001




Note: P/E was calculated as the current price divided by the average of earnings for the past 10 years adjusted for inflation.

Table 1. Real Stock Market Return in the Next 10 Years for Historical P/E Ranges of the S&P 500, 1927–2001 Data

P/E Range (low to high)	Median Return (annualized)	Worst Return (total)
5.6 to 10.0	11.0 %	46.1 %
10.0 to 11.7	10.6	37.3
11.7 to 14.1	10.0	4.1
14.1 to 16.7	9.0	–19.9
16.7 to 19.4	5.4	–23.1
19.4 to 32.6	–0.4	–35.5
32.6 to 45.0	Here Be Dragons!	

Figure 2. S&P 500 Trailing 20-Year Real Earnings Growth, 1891–2001



Note: Earnings growth is annualized.

Table 2. Average 10-Year Real Earnings Growth, 1927–2001 Data

Starting P/E	Retention Rate (%)				
	Negative to 37.7	37.7 to 44.4	44.4 to 50.3	50.3 to 63.9	63.9 →
5.9 to 10.4	4.1 %	2.5 %	2.2 %	–0.3 %	
10.4 to 13.8	4.3	2.5	2.4	0.6	
13.8 to 17.2	3.3	2.5	1.7	–0.4	
17.2 to 26.3	4.3	2.7	0.8	–0.6	
26.3 →					The Dragons Are Here!

has been weak. Some would argue that in an efficient market, the current P/E simply *must* be justified by high earnings expectations. Asness thinks the data do not provide much support for this proposition.

A third reason might be that high cash retention leads to above-normal growth. But referring to Table 2, he pointed out that the current retention rate has been significant in relation to real earnings growth and the retention at the beginning of a 10-year period is *inversely* related to the subsequent 10-year growth

in earnings! Why should this be? One answer is empire building. Retention of earnings is simply not productive. A second is a desire on the part of managers to smooth dividends. In any case, the current retention rate is about 65 percent, and Table 2 is not encouraging for the future of the stock market.

A second way in which the market might be saved is through low interest rates. Can low interest rates save stocks? Panel A of **Figure 3** is encouraging: Interest rates below about 3 percent are very helpful.


But Panel B shows that over a longer historical period, the news is not so good. The indicator seems to be the earnings yield, E/P , less the bond yield, Y . There is evidence that nominal earnings growth is correlated with inflation. The P/E , however, is mostly a real entity, and comparing it with nominal bond yields cannot be expected to have much long-term forecasting power.

Finally, the willingness of investors to accept low stock returns might save the market. Are investors

willing to accept low stock returns? Declining volatility may be justifying high P/E s and low returns. **Figure 4** provides support for this idea, although the vertically plotted E/P minus Y mixes real and nominal data.


Figure 4 seems to work for the short term. The point on the graph for September 30, 2001, represents a high P/E coupled with a low ratio of realized 20-year stock-to-bond volatility. For the longer term, the E/P is a better guide to real stock returns.

Figure 3. S&P 500 "Absolute" and "Relative" Value



Note: S&P 500 P/E and E/P ; 10-year T-bond yield.

Figure 4. Stock versus Bond Valuation, 1950–2001




Note: S&P 500 E/P; 10-year T-bond yield.

Acceptance of a 6–7 percent nominal stock return appears not unreasonable. But Asness went on to present evidence that investors do not actually think they are facing such low returns. In this case, when they realize the true prospects, then short- to medium-term returns will be low. To raise the expected return on the S&P 500 by 2 percentage points, the price must fall about 50 percent.

Figure 5 shows the results of forming long–short portfolios (based on Wall Street growth forecasts) in which the portfolios were long the high growers and short the low growers. The rolling 24-month beta of the portfolios has been consistently positive and, in recent years, has been massively positive. Every rally has seen the high-expected-growth stocks crushing the low-expected-growth stocks. Asness thought this

Figure 5. Rolling 24-Month Beta of Long–Short Portfolio, December 1983–September 2001



Note: Except for 2001, dates are as of December.

was not a picture of investors willing to accept lower equity premiums.

In conclusion, he said:

- Broad stock market prices are still well above the levels of most recorded history (and of all history excluding 1999–2000 and just before the crash of 1929). Unless a miracle happens, we must prepare for very low returns as compared with history.

- The choice is between low long-term returns forever and very low (crash type) returns followed by more historically normal returns.

Finally, he offered the following reflection: Do the events of 1999–2001 strike anyone as a picture of rational investors accepting a permanently low risk premium? Answer: No.

Theoretical Foundations: Discussion

Stephen Ross (*Moderator*)

Robert Arnott

Clifford Asness

Ravi Bansal

John Campbell

Bradford Cornell

William Goetzmann

Roger Ibbotson

Martin Leibowitz

Rajnish Mehra

Thomas Philips

Robert Shiller

Richard Thaler

STEPHEN ROSS (*Moderator*)

I have a few brief comments. They will be brief for two reasons. First, I am confused. Second, even in my confusion, I am in the uncommon position of not having a lot to say. Let me turn first to Cliff Asness's presentation.

What is puzzling to me about Cliff's presentation is that the discussions about P/Es and other broad descriptors of the market seem to me to be discussions that we could have held 100 years ago. The vocabulary would have been a little different, but in fact, not only could we have held the discussion, I suspect these discussions *were* held 100 years ago. So, I don't think we are saying many things differently now than we said back then.

What is troubling to me is that we are supposed to be making progress in the theory. To the contrary, the theory seems to me to be in a wasteland, not just regarding the risk premium but, more generally, in much of finance. We are in a period of time, a phase, in which data and empirical results are just outrunning our ability to explain them from a theoretical perspective. This position is a very tough one for a theorist who used to dine high on the hog when we had derivatives pricing, where theory worked wonderfully. Now, we are interested in theory to explain the problems, which is not working quite so wonderfully.

It seems to me that the issues involving P/Es are issues involving whether or not these processes are mean reverting. Obviously, something like the P/E

has to revert to the mean; it is only a yield. Jonathan Ingersoll made a wonderful comment about interest rates and whether interest rates revert or not. He noted that interest rates existed 4,000 years ago in Egypt and if interest rates didn't mean-revert, they would be 11,000 percent today. So, they have to revert.

We know P/Es revert, but they seem to revert very slowly, and we are able to measure the reversion only with great difficulty. Our efforts to measure, for example, stock returns—not actual returns but expected returns—have basically been futile.

I also have some comments about Richard Thaler's presentation. I am often characterized as a defender of the neoclassical faith. I know I am because often I am asked to debate Richard. Sometimes, however, I am characterized as a shill of the neoclassical school. So, it is not clear to me which position I am supposed to represent in the minds of market pundits. But I will say that I feel a bit like one of those physicians with a gravely ill patient to whom I would like to suggest the possible benefits of herbs and acupuncture—alternative medicine. I call for "alternative finance," not behavioral finance as the alternative approach, but an alternative that may offer a little bit of hope.

What I actually think is that our prey, called the equity risk premium, is extremely elusive. We cannot observe the expected return on stocks even with stationarity in time-series data because volatility and the short periods of time we are able to analyze give us little hope of actually pinning down a result. The best hope, from the empirical perspective, seems to lie in cross-sectional analysis, which is not what we are talking about here; we are talking mostly about *time series*, for which we do not have many observations. Cross-sectional analysis says that the excess returns should be the risk premium times the beta. If we could find some way to spread excess returns, maybe through P/Es of individual stocks, then we'd have a better chance of measuring expected return at each point in time—no matter what theory we decide to pin our hopes on.

The theory itself is a myth, and in this case, Richard and I are in complete agreement. Any hope of tickling, or torturing, some reasonable measure of the risk premium out of consumption data is forlorn. It resides in the hope that somehow people are rational.

I love old studies. For example, in one study on consumption data that was done mostly in Holland, the researchers observed shoppers in supermarkets

to see what happened when the price of soap was higher than the price of bread. These shoppers did not adjust their marginal rates of substitution to the prices of consumer goods at a single point in time, let alone in the presence of uncertainty and over time. But consumption theory has always said that people would adjust their marginal rates of substitution for prices that evolve over time in a stochastic world.

I am not at all surprised, nor am I troubled, by the fact that we do not find any meaningful correlations between something that we may or may not be able to measure, such as expected return and consumption, and the interplay between them. So, I applaud Richard's view that we ought to consider other reasons to explain why people do what they do.

The real puzzle may be: Why do investors behave the way they do based on what the premiums actually are? And here too, I have to say that even though neoclassical theory is not up to the task of explaining this behavior, and it is not doing a good job, I am not sure that behavioral theory has much more to say to us.

Behavioral anecdotes and observations are intriguing. Behavioral survey work is empirically fortified. But behavioral theory does not seem to have a lot of content yet. In interpreting the study that Richard mentioned about the incompatibility of two gambles, one has to be very careful. Those gambles are incompatible if they are assumed to hold over the entire range of the preference structure. But there is no reason to believe that the gamble holds over the entire range of the preference structure. We do not believe that if the guy wins \$20 million he won't take the 110 to 100 gamble. The uniformity requirements in that assumption bend the question. A lot of curious things are going on in those kinds of analyses of behavioral assumptions. And even the richer models, such as those of DeLong and Shleifer (1990), have their own problems.

In summary, I am a theorist and I am confused. I would like theory to make progress, and I would like for us to be able to address some of these issues successfully. I do not really care whether we do so from a neoclassical or another perspective, but I find myself facing an enormous, complicated array of phenomena that come under the heading of "the equity risk premium puzzle" and I'm completely unable to explain any of it.

RAJNISH MEHRA: One thing that Richard Thaler missed was that most of these models do not incorporate labor income. Constantinides, Donaldson, and I (1998) have been doing work in this area for the last couple of years. We have been analyzing the implications of the changes in the characteristics of labor income over the life cycle for asset pricing. The

idea is simple: The attractiveness of equity as an asset depends on the correlation between consumption and equity income, and as the correlation of equity income with consumption *changes* over the life cycle of an individual, so does the attractiveness of equity as an asset. Consumption can be decomposed into the sum of wages and equity income. A young person looking forward in his or her life has uncertain future wage *and* equity income; furthermore, the correlation of equity income with consumption will not be particularly high as long as stock income and wage income are not highly correlated. This is empirically the case. Equity will thus be a hedge against fluctuations in wages and a "desirable" asset to hold as far as the young are concerned.

Equity has a very different characteristic for the middle-aged. Their wage uncertainty has largely been resolved. Their future retirement wage income is either zero or fixed, and the fluctuations in their consumption occur from fluctuations in equity income. At this stage of the life cycle, equity income is highly correlated with consumption. Consumption is high when equity income is high, and equity is no longer a hedge against fluctuations in consumption; hence, for this group, equity requires a higher rate of return. The way Constantinides, Donaldson, and I approach this issue is as follows: We model an economy as consisting of three overlapping generations—the young, the middle-aged, and the old—where each cohort, by the members' consumption and investment decisions, affect the demand for, and thus the prices of, assets in the economy. We argue that the young, who should be holding equity, are effectively shut out of this market because of borrowing constraints. In the presence of borrowing constraints, equity is thus exclusively priced by the middle-aged investors, and we observe a high equity premium. We show that if there were no constraints on young people participating in the equity markets, the equity premium would be small.

So, I feel that life-cycle issues are crucial to any discussion of the equity premium.

JOHN CAMPBELL: I want to follow up on the point Rajnish Mehra made because one part of Richard Thaler's talk was normative analysis—the claim that if the equity risk premium is as much as 4–5 percent, long-term investors should obviously hold their money in stocks or even leverage a position to hold their money in stocks. I think that, as a normative statement, that prescription is simply wrong.

I am going to take as a benchmark a model with constant relative risk aversion at some reasonable, traditional low number. The simple formula for the share you should put into stocks if you are living off

your financial wealth alone and if returns are distributed identically every period is as follows: the risk premium divided by risk aversion times variance. Suppose the risk premium is 4 percent and the standard deviation of stocks is 20 percent; square that and you get 4 percent. Now, you have 4 percent divided by risk aversion times 4 percent. So, if your risk aversion is anything above 1—say, 3 or 4—you should be putting a third of your money in stocks or a quarter of your money in stocks. It is just not true that with low risk aversion and a risk premium of 4–5 percent you should put all your money in stocks.

So, what's happened to the puzzle? Why don't I get an equity risk premium puzzle when I look at it from this point of view? Well, the key assumption I made is that *you are living off your financial wealth entirely*. It follows then that your consumption is going to be volatile because it will be driven by the returns on your financial wealth. The only way to get an equity risk premium puzzle is that when you look at the smoothness of consumption, you see that it is much smoother than the returns on the wealth portfolio. Why is that?

Rajnish's point is that other components of wealth, such as human capital, are smoother, which is keeping down the total risk of one's position. If you have these other, much smoother human assets, then of course, stocks look very attractive. But I think it's important not to assert that a risk premium of 4 percent should induce aggressive equity investment.

I am reminded of Paul Samuelson's crusade over many years to get people to use utility theory seriously, as a normative concept. He was always trying to combat the view that you should just maximize the expected growth rate of wealth. He got so frustrated by his inability to convince people of this that he finally wrote an article called, "Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long" (1979). It is a wonderful article, and the last paragraph says, "No need to say more, I've made my point and but for the last word, I've done so in words of but one syllable." And every word in the article is a one-syllable word except for the last word. It is almost impossible to read, of course, but the point is important: We may not want to use standard utility theory as a positive theory, but we should try to use it as a normative theory, in my view.

ROSS: If you are going to use it as a normative theory, though, you do not have to place your attention entirely on the constant relative-risk-aversion utility function. The broader class of linear risk-tolerance models has exactly the same function (with the addition of deterministic parts to the income stream), except they work in the opposite direction.

So, if someone has a linear risk tolerance with a high threshold for that risk tolerance, then the equity risk premium puzzle reappears because the desire to invest is huge even when the risk premium is relatively low.

RICHARD THALER: Let me respond briefly. You have all these models that are based on consumption, and it is true (and I appreciate John Campbell's clarification) that to really understand this puzzle, you need to emphasize consumption smoothing. Otherwise, you get precisely the result that John suggested.

But the puzzle I was informally identifying before refers to other investors that I think have been neglected in much of this theoretical research. Those simulations that Marty Leibowitz was doing were mostly for defined-benefit pension funds, and I did some similar simulations for a foundation that I've been associated with over the years. Foundations have 5 percent mandatory spending rules. Now, if you crunch the numbers and you are investing in bonds, basically you are certain to be out of business in the near future unless you can find some bonds providing a 5 percent real rate of return. With TIPS we were getting close for a while.¹ But if the real interest rate is 2 percent and you have to spend 5 percent, you are soon going to be out of business. One question I have for the theorists, of which I am not one, is: What's the normative model we want to apply for those investors and what does it tell us about the kind of risk premium we should expect?

BRADFORD CORNELL: I have one question: Most of you are involved in one way or another with investment firms, and it is almost a mystery to me that you read academic papers where you see things like "consumption process," "labor income," "risk aversion," and so on, and then you attend an actual investment meeting—where none of these concepts are even remotely talked about. So, how do you bridge the gap between the supposed driving factors of the models and equilibrium returns and the way people who are actually making decisions make them? Is there a way to tie all of it together?

ROSS: There does seem to be a disconnect between the two areas and the two literatures. It is, actually, a fundamental theoretical disconnect. In these markets, with their many institutional players, the institutions are typically run by managers under some type of agency structure. So, there must be some sort of agency model for the people who run the pension funds and other institutions. They are the ones who

¹ Treasury Inflation-Protected Securities; these securities are now called Treasury Inflation-Indexed Securities.

make investment decisions. In the theoretical structures we build that include consumption, we seem to have the view, or maybe just the wishful thinking, that whatever the underlying forces in the economy are, these institutions will simply be transparent intermediaries of those forces, so the agents who are representing these institutions will simply be players in people's desire to allocate consumption across time or will be dealing with the life-cycle problems of people. Some take a Modigliani view that the *people* will adjust their actions around whatever the agents do. The net result is that the actions of the agents and the people coincide, which seems to me overly hopeful. I don't believe it is the case.

CLIFFORD ASNESS: Is it more complicated than saying the description Richard Thaler presented works better for what actually happens in a boardroom than any of the theory? Behavior like myopic loss aversion is true. Many of us have behaved that way. The fact that people make choices in the ways that they do does not have to be proven by a survey. As a manager who has gotten way too much money after a good year and too many redemptions after a bad year, I can tell you people focus on the short term.

I have one comment about Steve Ross's initial response. I don't think anyone would argue about the fact that P/Es are mean reverting. But that is not the exciting part of the puzzle. The exciting part, which is incredibly challenging, is that if we all accept that P/Es are mean reverting to an unconditional mean, what we are disagreeing about is what that unconditional mean either should be, in theory, or is. Mean reversion is a pull toward something, and the open issue is not mean reversion but whether the "right" (meaning unconditional mean) P/E is 15. If it is and we are in the high 20s, then mean reversion is not going to work as a good model for the next year. But the pull was downward for a long time, so I do not think my comments were trying to be insightful about P/Es being mean reverting. They have to be, or else they are unbounded in some direction.

MARTIN LEIBOWITZ: This is just strictly an observational comment, not a theoretical one, and it has to do with the comment about myopic loss aversion or myopic return attraction, which is the other side of the coin. As Cliff Asness said, there's clearly some pain in the short term and also some joy in the short term, depending on your outcomes. But I think what actually happens is that people incorporate a kind of Bayesian revision, that the prospects for the future are based on what have been the most immediate

short-term returns.² We see it in terms of the flow of funds into, for example, TIPS—a wonderful instrument with a great yield, a +4 percent real rate. We couldn't get anyone to invest in them until, suddenly, we had a 12.76 percent return year in the equity market, at which point, of course, the real return on equities was a lot lower than it had been and money started flowing into TIPS big time. Short-term return is a very powerful force.

THALER: Aren't you too Bayesian, then, to be sarcastic?

LEIBOWITZ: Yes, Bayes would recoil because in the fixed-income area, this short-term focus is clearly, you know, a kind of nuttiness, although there's something to it. It does show that real rates can decline. I think some people were thinking: Why were we stuck with real rates in the area of +4 percent? So, myopic loss aversion is not totally irrational, even in the fixed-income area. In the equity area, where the risk premium is so elusive and unmeasurable, I think that investors do place a lot of weight on these myopic results, and not just in the short term; they are interested in what the data say about the long term.

ASNESS: Can we call it Bayesian without priors?

LEIBOWITZ: I think there are priors. I think there really is a Bayesian division going on.

THALER: I want to explain that in the study by Marty Leibowitz, which I so meanly presented, one of the conclusions he reached is that those 20-year numbers look really, really good but that the plan sponsors, the target audience of Marty's study, were going to have to answer some difficult questions over the next two or three years. This problem is an agency problem. The investment committee or whoever is making the investment decisions will get a lot of heat if lots of losses occur on their watch. Typically, the manager running the pension plan is going to be in that job for only two or three years and will then rotate into another job.

ROSS: That agency problem exacerbates this issue even further. With the distinction between the real economy (represented by Rajnish Mehra and John Campbell) and the financial markets, the transmission

² Bayes' Law determines a conditional probability (for example, the probability that a person is in a certain occupation conditional on some information about that person's personality) in terms of other probabilities, including the base-rate (prior) probabilities (for example, the unconditional probability that a person is in an occupation and the unconditional probability that the person has a certain personality).

mechanism through institutions becomes even more difficult to explain. Are those who run institutions subject to a variety of psychological vagaries of this sort? Why, if this is an agency problem, has it been so poorly solved to date? It seems to throw up even more theoretical puzzles for us.

LEIBOWITZ: Just a real quick response. That research of mine that Dick Thaler mentioned actually spurred a whole series of papers in which we looked at all kinds of reasons why people would not be 100 percent in stocks. We looked at it from all kinds of different angles—both theoretical and empirical—and we always kept getting this kind of lognormal type of distribution with nice, beautiful tails; it was pretty weird never to see underperformance over long periods of time.

The only conclusion we could finally come to was that, basically, as people peer into the future, they see risk. They are not talking about something with volatility characteristics. They are not talking about return that behaves in a linear fashion. But they see something out there that, basically, fundamentally, scares them. They can't articulate it, but it keeps them from being 100 percent in stocks.

CAMPBELL: I want to defend the relevance of consumption, even in a world with both behavioral biases and agency problems. It would be ludicrous to deny the importance of those phenomena, but even in a world with those phenomena playing a major role, consumption should have a central role in our thinking about risk in financial markets. In the long run, consumption drives the standard of living, which matters to people. So, consumption is a very influential force in investors' decisions.

Can consumption models be applied to endowments, to long-term institutions? I argue that they can, and I have some knowledge of this issue from talking to the managers of the Harvard endowment. Harvard's new president, Lawrence Summers, is trying to make sense of Harvard's spending decisions, which have always been made on an *ad hoc* basis. The endowment maintains very stable spending for a number of years, and then spending rises periodically. Now, in many universities, endowments generally have a smoothed spending rule, so spending levels are linked to past spending levels and the recent performance of the endowment. This rule makes perfect sense if you think that universities get utility from spending but also have some sort of habit formation. It is internal as related to their own history: They hate to cut the budget because it is really painful, the faculty are up in arms, and the students are

screaming. And it is related to external situations: They hate to fall behind their competitors. I know that the Harvard endowment managers look very carefully at the management of the Yale endowment, because there's nothing worse than having Yale outperform Harvard. So, habit formation and consumption spending are extremely relevant to endowments. The relationship may be a little more complicated than just saying, "Oh, they have power utility," but you can make sense of the way they think by reference to spending, not only at the micro level but also in terms of the aggregate consumption in the economy.

In the long term, the correlation between consumption growth and the stock market has been quite strong—in the United States and in other countries. And it makes sense. We know that when the economy does well, the stock market does well, and vice versa. There is a link, a correlation, and it represents a form of risk over the longer run.

Aggregate consumption is also an amazingly accurate measure of the sustainable long-term position of the economy. We know that consumption, financial wealth, and labor income are all held together by budget constraints. You can't let your consumption grow indefinitely without some reference to the resources that are available to support it. So, no matter what the behavioral influence is, there is still a budget constraint that is bound to hold consumption, wealth, and income together. You can ask the empirical question when you look at the data: What adjusts to what? If you have a behaviorist's view, you might think that consumption would adjust to the harsh realities of the budget constraint over time. Instead, what seems to happen is that consumption follows a random walk—as if it is set to the level that is sustainable at each point in time. When wealth gets out of line or income gets out of line, they adjust to consumption. So, there's short-term volatility in the financial markets, but when financial wealth is very high relative to consumption, what tends to happen is financial wealth falls. That is just a fact, it does not suggest a particular model, but I think it does suggest the relevance of consumption—together with agency problems and very interesting and important behavioral phenomena—in thinking about the markets.

CORNELL: If consumption is relevant, what type of information would you expect to see flowing through the pipeline of an organization such as TIAA-CREF? How would you expect to see information flowing from the ultimate clients, who are the consumers, into the organization so that the organization can act as the agent on their behalf?

CAMPBELL: Well, TIAA-CREF is running a defined-contribution pension plan. So that, in a sense, information does not have to flow into it. But it seems to me the way to think about defined-benefit pension plans is that they have evolved over a long period of time to reflect the conservatism of the ultimate clients. For example, labor unions negotiate pension arrangements to give their members very stable income in retirement. And even if we accept that agency problems introduce imperfections, it seems to me that the liabilities defined-benefit pension plans have are very stable because of an expressed preference for stable consumption streams.

THALER: The residual claimant to those plans is the company, and the company is supposed to be virtually risk neutral. So, I think the model John Campbell described, which is sort of a habit-formation model, has some plausibility to it as applied to endowments. What is more difficult is to try to use that model in explaining the behavior of the typical plan sponsor of a defined-benefit pension plan.

ROBERT SHILLER: The general public of investors does not, of course, have an economic model like those produced by economists. They do, however, know the definition of stocks and bonds. They know that bondholders get paid first and stockholders are the residual claimants after the bondholders are paid. They know that. The original idea for a stock market was that stockholders are the people who can bear risk and that buying stocks is designed to be a risky contract—which, I think, is very much on investors' minds. So, if we tell them, "Well, in this last century, we were really lucky. Nothing really went wrong. We had five consecutive 20-year periods in which stockholders did really well," I believe that investors then think, rationally, that what we are telling them about low risk for stocks is pretty unconvincing. Investing in stocks is still investing in an asset that was designed for people who can take a lot of risk. There are no promises, and the government isn't going to bail you out if the stock market collapses. The government is perfectly free to throw on a big corporate profits tax; they've moved it up and down. And the shareholder gets no sympathy when the government does so. So, people are rational to be wary, to require a high expected return to take that risk.

ROBERT ARNOTT: I think in this whole discussion of risk premiums we have to be very careful of definitions. In terms of expected returns on stock, there is the huge gap between rational expectation based on a rational evaluation of the sources of return, current market levels, and so forth, versus hope. The inves-

tors out there are not investing because they expect to earn TIPS plus 1 percentage point.

And we have a semantic or definitional problem in terms of past *observed* risk premiums, exemplified by the Ibbotson data, between a *normal or unconditional* risk premium, which a lot of the discussion so far seems to have centered on, and the *conditional* risk premium based on current prospects. So, one of the things that we have to be very careful of is that we clarify what we're talking about—past observed risk premiums, normal (unconditional) risk premiums, or conditional premiums based on current prospects.

ROGER IBBOTSON: We have talked mostly about either the behavioral perspective or the classical (or neoclassical) perspective. The classical approach can be interpreted or reinterpreted in many ways as we get more and more sophisticated in our understanding of what the risk aversion might be for the predominant people in the market. And we can put behavioral overlays on classical theory. Ultimately, I think this topic is a rich land for research, and I encourage it, but we are not very close now to getting a fix on an estimate for the risk premium. At first, it appeared that theory suggested low risk premiums, as per Mehra and Prescott (1985), but I think at this stage of the game, using classical theory with behavioral overlays, we can't pinpoint the answer.

THOMAS PHILIPS: An idea that ties together many of the discussions associated with the risk premium is the notion of how to estimate something if you don't have a model or if you're not sure what you are doing. The typical answer is to take the historical average or the sample mean. If we stop to consider why investors buy TIPS at certain times and pull out of hedge funds at other times, we find, more often than not, that the answer is grounded in their use (and abuse!) of the sample mean of the historical returns of that asset class. The trouble is that the sample mean is a terrible estimator. It is easy to show that the sample mean can have huge biases; you just have to vary the risk premium a little bit, for example, or have slightly different economic assumptions, and the estimate and reality diverge sharply. But the sample mean does seem to be the driving force behind most people's behavior. What you observe at cocktail parties or working with clients is this enormous drive toward investing in the asset class with the highest historical return. And I believe it is a fundamentally bad way to think about the problem.

MEHRA: I want to say a couple of things in defense of neoclassical economics. First, for psychological vagaries and other behavioral phenomena to affect prices, the effect has to be systematic. Unless these

phenomena occur in a systematic way, the behavior will not show up in prices. So, one has to be very careful about saying, "This is how I behave so I should model market behavior that way." Many of our idiosyncrasies may well cancel out in the aggregate.

Second, most of our economic intuition is actually based on neoclassical models. Ideally, new paradigms must meet the criteria of cross-model verification. Not only must the model be more useful for organizing and interpreting observations under consideration, but it must not be grossly inconsistent with other observations in growth theory, business cycle theory, labor market behavior, and so on. So, I think we should guard against this tendency of model proliferation in which one postulates a new model to explain each phenomenon without regard to cross-model verification. A model that is going to explain one part of reality but then is completely inconsistent with everything else does not make much progress. That is my biggest concern.

ROSS: It seems to me also that there is a vocabulary issue at work here. We have heard the phrase "habit formation" used by many people to mean many different things. On the one hand, the term is used by the behavioralists as though it is some kind of psychological phenomenon. On the other hand, John Campbell uses it as a description of the way universities behave. In either case, it is difficult to tell the difference between whether some fundamental underlying costs that universities face produce a behavioral pattern that looks like habit formation on the preference side but might have nothing to do with it or whether the universities' preferences are perfectly independent across time, are intertemporally independent, but the basic cost structure induces a net behavior that looks like they're concerned about what they did in the past or they are concerned about preserving what they did in the past.

The same is true on the behavioral side. It could well be that there is some fundamental psychological underpinning that we can argue for in terms of habit formation. All you are really saying is that, on the preference side, people don't have adequately separable preferences all the time, that there is some induced link between preferences at one point in time and consumption at one point in time and consumption at another time. There may be some substitutability that we are not capturing in the additive case. So, I think that all of these phenomena have the funny and interesting property that both the neoclassical economist and a purely psychological economist, or behavioral economist (I don't know what the proper phrase is anymore), could wind up saying that the reduced

form could be the same for both of them. They just have different ways of getting there.

SHILLER: I think the difference between behavioral economics and classical economics is totally a difference of emphasis. The behaviorists are more willing to look at experimental evidence, a broad array of evidence. Indeed, expected utility is a behavioral model; psychologists also talk about expected utility. So, I think the difference is somewhat methodological; it is not a subject matter difference. It is a question of how willing you are to experiment with different variations.

THALER: Well, habit formation is obviously to some extent a description of preferences. Nothing says it's irrational. The simple additive (and separable) model is the easiest to use, so we naturally started with that model. But you could add completely hypo-rational agents who have preferences that change from one period to another, and you could, of course, have agents who are making the so-called Bayesian forecasts that Marty Leibowitz referred to with those same preferences.

ROSS: There are some exceptions, though, like framing or path dependence. Those tend to be time inconsistent, and time consistency is required in what we typically think of as rational models.

WILLIAM GOETZMANN: A lot of interesting theoretical work is going on, but I want to put in a plug for empirics. Theorists have looked at the price behavior of markets and of individual securities, but a lot of the models have this behavioral component, rational or otherwise, at their heart—whether in identifying the marginal investor or what have you. Yet, we have almost no information about how actual investors behave. Organizations have a lot of that information, but it may never see the light of day for our research purposes. We're beginning to see a little bit of this information cropping up here and there (and sometimes companies that allow us to have it are sorry they did). But imagine the ability to take hundreds of thousands of accounts, time series of accounts, identify the people who seem to exhibit myopic loss aversion, and then test to see whether their behavior has any influence on prices. That work would provide a way to identify whether pathologically behaved people have a short-term or a long-term influence on price behavior. In the long run, empirical study is how we are going to be able to answer some of these questions.

RAVI BANSAL: There is a lot of discussion about preferences, and many of the implementations of this theory lead to the result that asset price fluctuations are a result of cost-of-capital fluctuations. The models do not have much room for expected growth rates. The models build on a long-held belief in economics that consumption growth rates and dividend growth rates are very close to being identically and independently distributed (i.i.d.). It is the notion that most people have. I think we need to rethink that idea. A lot of hidden persistent components are in these growth processes; the realized growth process looks like an i.i.d. process, but if these growth rates have a small persistent component, the ramifications are huge. Small persistent components of any of these growth rates would have dramatic implications for how we think about what is causing asset prices to fluctuate. Statistically, there is actually some evidence to support the view that there are some persistent components in both consumption and growth rates. If such components are put into a model, the unforeseen components can explain equity premiums because consumption goes up at the same time dividends go up. News about consumption and dividend growth rates continuously affects perceptions about long-run

expected growth rates, which leads to a lot of asset volatility. This channel is important for interpreting what goes on in asset markets.

Behavior is important, clearly, but understanding the dynamics of cash flows, of consumption, is equally, if not more, important. So, in a paper that Amir Yaron and I wrote (Bansal and Yaron 2000), we allowed for that possibility. And we actually show that when you rely on the Epstein–Zin (1989) preference structure and allow for intertemporal elasticity of substitution to be more than 1.0 (which makes intuitive sense to me), then you can actually get the result that during periods of high anticipated consumption growth rates, the wealth-to-consumption ratio rises. So, in terms of the asset markets, asset valuations will rise simply because of higher expected growth rates. When you require the intertemporal elasticity of substitution to be more than 1.0, then when people expect good times, they want to buy assets. I find this quite intuitive. When you allow for this possibility, you can explain through these neoclassical paradigms a lot of the equity premium and volatility in the market. So, focusing on aggregate output growth is a pretty important dimension.

Historical Results I

Jeremy J. Siegel

*Wharton School of Business
Philadelphia*

Analysis of the very long term in U.S. markets indicates that average real stock market returns have been about 7 percent and average real T-bond and T-bill returns have been about half that figure. Downward bias in the more recent bond returns and upward bias in recent valuations may be skewing the analysis. Valuations have been rising for three possible reasons: declining transaction costs, declining economic risks, and investors learning that stocks have been undervalued on average throughout history. An analysis of the historical relationships among real stock returns, P/Es, earnings growth, and dividend yields and an awareness of the biases justify a future P/E of 20 to 25, an economic growth rate of 3 percent, expected real returns for equities of 4.5–5.5 percent, and an equity risk premium of 2 percent (200 bps).

Table 1 shows historical returns and the equity risk premium (on a compounded and an arithmetic basis) for the U.S. markets from 1802 through September 30, 2001. The last columns display the equity risk premium based on a comparison with U.S. T-bonds and T-bills, which is just the difference between the real return for stocks and the real return for bonds and bills. I broke out these returns and premiums into the three major sub-

periods since 1802 and also into 20-year post-World War II periods.

When I wrote the book *Stocks for the Long Run* (Siegel 1998), I was struck by the fact that for all the very long periods (and the definition of “long” is more than 50 years), the average real annual stock market return is just about 7 percent a year, maybe a tad under. This return also holds true for the three sub-periods 1802–1870, 1871–1925, and 1926–2001 and for the whole 1946–2001 post-WWII period. (By the way, almost all of the inflation the United States has suffered over the past 200 years has come since World War II, and as we economists should not find surprising, stocks—being real assets—were not at all adversely affected by post-WWII inflation). So, 7 percent appears to be a robust measure of the long-term annual real stock return.

For periods of several decades, however, the real return on stocks can deviate quite a bit from that 7 percent average. Some of those extreme periods since WWII include the bull market of 1946–1965, the bear market of 1966–1981, and the great bull market that lasted from 1982 to the end of 1999. From 1982 through 1999, the average real return on stocks was 13.6 percent, which is double the 200-year average.

That recent experience may color investors’ estimates of the equity risk premium today. In the roundtable Discussion for the opening session [“Theoretical Foundations”], there was talk about Bayesian updating, and I do believe that investors place greater weight on the more recent past than we economists think they should. Perhaps investors believe that the underlying parameters of the system have shifted or the model or paradigm has changed or whatever, but I think some of the high expectations investors have for future returns have certainly come from the recent bull market. For many investors, their bull market experience is the only experience they have ever had with the markets, which could certainly pose a problem in the future if excess-return expectations are widespread and those expectations are frustrated.

Table 1. Historical Returns and Equity Premiums, 1802–September 2001

Period	Real Return						Stock Excess Return over			
	Stocks		Bonds		Bills		Bonds		Bills	
	Comp.	Arith.	Comp.	Arith.	Comp.	Arith.	Comp.	Arith.	Comp.	Arith.
1802–2001	6.8 %	8.4 %	3.5 %	3.9 %	2.9 %	3.1 %	3.4 %	4.5 %	3.9 %	5.3 %
1871–2001	6.8	8.5	2.8	3.2	1.7	1.8	3.9	5.3	5.0	6.6
<i>Major subperiods</i>										
1802–1870	7.0 %	8.3 %	4.8 %	5.1 %	5.1 %	5.4 %	2.2 %	3.2 %	1.9 %	2.9 %
1871–1925	6.6	7.9	3.7	3.9	3.2	3.3	2.9	4.0	3.5	4.7
1926–2001	6.9	8.9	2.2	2.7	0.7	0.8	4.7	6.2	6.1	8.0
<i>Post World War II</i>										
1946–2001	7.0 %	8.5 %	1.3 %	1.9 %	0.6 %	0.7 %	5.7 %	6.6 %	6.4 %	7.8 %
1946–1965	10.0	11.4	–1.2	–1.0	–0.8	–0.7	11.2	12.3	10.9	12.1
1966–1981	–0.4	1.4	–4.2	–3.9	–0.2	–0.1	3.8	5.2	–0.2	1.5
1982–1999	13.6	14.3	8.4	9.3	2.9	2.9	5.2	5.0	10.7	11.4
1982–2001	10.2	11.2	8.5	9.4	2.8	2.8	1.7	1.9	7.4	8.4

Note: Comp. = compound; Arith. = arithmetic.

Sources: Data for 1802–1871 are from Schwert (1990); data for 1871–1925 are from Cowles (1938); data for 1926–2001 are from the CRSP capitalization-weighted indexes of all NYSE, Amex, and Nasdaq stocks. Data through 2001 can be found in Siegel (2002).

The annual real bond returns provided in Table 1 show an interesting trend. From 1802 through September 30, 2001, the average annual real T-bond return was 3.5 percent, about half the equity return. In the major subperiods, this return has been trending decidedly downward. Beginning in the 19th century, it was nearly 5 percent; it then fell to 3.7 percent in the 1871–1925 period; it was 2.2 percent for the 1926–2001 period; and since the end of WWII, it has been only 1.3 percent. From 1982 onward, as interest rates and inflation have fallen, bonds have produced a much greater real return than average. When I was studying finance in the 1970s, we learned that both T-bill and T-bond real returns were close to zero. Yet, over the past 20 years, those real returns have definitely risen.

When TIPS were first issued, they were priced to yield a real return of 3.5 percent, which is close to the average 200-year long-term real return of bonds.¹ Investors rightfully ignored the low real returns on bonds of the past 75 years (the period made popular by Ibbotson and the standard benchmark for the profession) in determining the TIPS yield. In fact, in 2000, during the stock market boom, TIPS were priced to yield a real return of almost 4.5 percent. Currently, the long-term TIPS yields have fallen back to a 3.0–3.2 percent range, depending on the maturity.

The real returns on T-bills tell the same story as for bonds, although for bills, the return is generally a bit lower. Of course, bills do not generate the capital

gains and losses that bonds do, so in the post-WWII period, bill returns have not fluctuated as much as bonds. Note that from 1982 forward, the annual real return for bills is 2.8 percent, far higher than the nearly zero average real return realized in the previous 55 years. In other words, periods as long as a half century can be quite misleading in terms of predicting future returns.

The problem is that while real stock returns were maintaining their long-term historical average real return of about 7 percent, real bond and bill returns were very low over the past 75 years, particularly up to 1980. Recognition of this phenomenon might help us understand why the equity premium has been so high in data from 1926 to the present.

The equity premium calculated for the past 75 years is biased downward for two reasons—bias in bond returns and bias in equity valuations.

Bias in Bond Returns

First, real historical government bond returns were biased downward over the 1926–2001 period. I say so because all the evidence points to the fact that bondholders simply did not anticipate the inflation of the late 1960s and 1970s. Investors would not have been buying corporate and government bonds of 30-year duration with 3.5 percent coupons (as they did in the 1960s) had they had any inkling of the inflation risk. I attribute part of that ill-fated confidence to the fact that few had a complete understanding of the inflationary implications of the shift from a gold-based to a paper monetary standard.

¹ TIPS are Treasury Inflation-Protected Securities; these securities are now called Treasury Inflation-Indexed Securities.

The gold standard was prevalent during the 19th century and much of the early 20th century when prices were stable over the long term. The United States (and most of the rest of the world) went off the gold standard in the early 1930s, but the effect was not immediately apparent. Although we had a pop of inflation following World War II, inflation was quite low up to the mid-1960s. So, in the 1960s, bond buyers were pricing 30-year bonds as if 30 years later their purchasing power would be nearly the same.

As inflation accelerated, bond buyers began to catch on. Bond yields rose, bond prices fell, and real bond returns were severely depressed. Table 1 shows that during the 15-year period from 1966 through 1981, the real return on bonds was a *negative* 4 percent. That period was long, and its effect is to bias downward the real return of bonds over the longer 1926–2001 period. I thus believe we should use higher real returns on fixed-income assets in our forecasting models, returns that are consistent with the real return on TIPS of 3–4 percent.

Bias in Equity Valuations


The second reason the equity risk premium is too high is that historical real stock returns are biased upward to some extent. **Figure 1** plots historical P/Es (defined here as current price of the S&P 500 Index divided by the last 12 months of *reported* earnings) from 1871 through September 2001. The straight line is the 130-

year mean for the P/E, 14.5. The latest P/E is about 37, surpassing the high that was reached in late 1999 and early 2000. So, the collapse of earnings that we have experienced this year has now sent the P/E to an all-time high.

Let me add a warning here: Part of the incredibly high P/E that we have now is a result of the huge losses in a few technology companies. For instance, JDS Uniphase Corporation wrote down its investments \$36 billion in the second quarter of 2001. The write-down was in reported earnings, not in operating earnings, and translates into a 5-point drop in the S&P 500 Index's valuation. So, approach these recent data on reported earnings with caution; \$36 billion from just one company's write-down has a huge impact on the market. Some of the technology issues are now essentially out-of-the-money options. When we compute numbers like the P/E of the market, we are adding together all the earnings of all the companies and dividing that into the market value. Because one company has big losses, it sells at option value, but another company with positive earnings can sell at a more normal valuation level. Adding these together might lead to upward biases in P/Es.

Nevertheless, there is no question that P/Es have risen in the past 10 years. If the market's P/E were to return to the historical (since 1871) average of 14.5 tomorrow, the annual real return on equities would fall 50 bps. And if the P/E had always remained at its

Figure 1. Historical Market P/E, 1871–2001



Note: Ending month for 2001 is September.

historical average level but the dividends paid had been reinvested, the annual real return on equities would be 115 bps lower than where it is today. The reason is that much of the real return on equities comes from the times when stock prices are very depressed and the reinvested dividends are able to buy many more shares, boosting stock returns. Much of the historically high returns on stocks has come when the market was extremely undervalued and cash flows were reinvested at favorable prices.

I believe there are several reasons for rising valuation ratios.

■ *Declining transaction costs.* One reason for rising valuations is the extensive decline in equity transaction costs. One-way transaction costs were more than 1 percent of the value of the transaction as late as 1975; costs are less than 0.2 percent today.² In the 19th and early 20th centuries, the (two-way) costs of maintaining a diversified portfolio could have been as high as 2 percent a year, whereas today indexed funds enable even small investors to be completely diversified at less than 0.2 percent a year.

■ *Declining risk.* Another reason for rising valuations may be declining levels of real economic risk as the U.S. economy has become more stable. The increased stability of labor income has enabled workers to accept a higher level of risk in their savings.

■ *Investor learning.* We cannot dismiss the fact that investors may have learned about the long-term risk and return characteristics of stocks. If investors have learned that stocks have been chronically undervalued on average, and in particular during recessions and crises, they will be less likely to let prices become undervalued, which leads to higher average valuations.

■ *Taxes.* Tax law has become increasingly favorable to equities. And low inflation, because the capital gains tax is not indexed, causes after-tax returns to rise. There has also been a proliferation of tax-deferred savings accounts, although it is not clear whether the taxable or tax-deferred investor sets stock prices at the margin.

²Charles Jones of Columbia University discussed declining transaction costs in "A Century of Stock Market Liquidity and Trading Costs" (2001).

Historical Growth Rates

As Table 2 shows, the real return on stocks has been 7 percent for the 1871–2001 period and is almost exactly the inverse of the P/E. If you divide this period into two subperiods—before World War II and after World War II—the real return for stocks remains roughly 7 percent but the dividend yield drops significantly from the first subperiod to the second, as does the payout ratio, and earnings growth rises.

In his presentation, Cliff Asness mentioned that he could not find in the data an increase in earnings growth when the payout ratio decreased [see "Theoretical Foundations" session]. But his findings are inconclusive because of the confusion between cyclical and long-term trends. In a recession, because dividends remain relatively constant as earnings plummet, payout ratios rise and earnings fall. In the subsequent economic recovery, earnings growth is higher and appears to follow a high dividend payout ratio. But this phenomenon is purely cyclical. Over long periods, a drop in the payout ratio and a drop in the dividend yield are matched almost one-to-one with an increased growth rate of real earnings. I find this relationship comforting because it is what finance theory tells us should happen over long periods of time.

Projecting Real Equity Returns

The link between the P/E and real returns is given by the following equation:

$$\text{Expected future real returns} = \frac{E}{P} + g \left(1 - \frac{RC}{MV} \right),$$

where

E/P = earnings yield, the inverse of the P/E

g = real growth

RC = replacement cost of capital

MV = market value of capital

RC/MV = book-to-market value, or 1/Tobin's q

I will call it the "Tom Philips equation" for projecting the real return of equity (Philips 1999). (I modified the formula somewhat.) According to this equation, if replacement cost does not equal market value, then the link between the P/E and future real returns must be modified. If Tobin's q is not 1, you have to correct

Table 2. Historical Growth Rates, 1871–September 2001

Period	Real Stock Return	Average P/E	Inverse of Average P/E	Real Earnings Growth	Dividend Yield	Real Dividend Growth	Real Capital Gains	Average Payout Ratio
1871–2001	7.06%	14.45	6.92%	1.27%	4.66%	1.09%	2.17%	62.24%
1871–1945	6.81	13.83	7.23	0.66	5.31	0.74	1.32	70.81
1946–2001	7.38	15.30	6.54	2.08	3.78	1.57	3.32	50.75

the earnings yield for the growth rate in the real economy to find expected future real returns. According to the equation, when the market value of equity exceeds the replacement cost of capital, as is the case today, the earnings yield *underestimates* future returns. The reason is that higher equity prices allow companies to fund capital expenditures by floating less equity, thereby reducing the dilution that this investment entails.

How much downward is the earnings yield biased? The Tobin's q on the latest data that I have is about 1.2. It was about 1.5, or even higher, in 2000. With long-run real growth at 3 percent, the last term, $g[1 - (RC/MV)]$, adds about 50 bps to the forecast of real return going forward. It added more in 2000 because Tobin's q was higher. So, if the P/E settles

down to 20 (and I believe that a future P/E should not be back at 14 or 15 but that a higher P/E is justified for the reasons I listed previously) and we emerge from the recession, then in terms of a long-term trend, E/P will be about 5 percent. Add the half a percentage point for the cheaper investment to maintain capital and you get a 5.5 percent expected real rate of return for equities. If the P/E is 25 in the future, with $1/25 = 4$ percent, adding the growth correction produces an expected real return for equities of 4.5 percent.

Keep in mind that TIPS are now priced to yield a real return of about 3 percent. So, because I believe that the long-run P/E in the market will settle between 20 and 25, the real future equity return is about 5 percent and the equity risk premium will be 2 percent (200 bps).

Historical Results I

Jeremy J. Siegel

*Wharton School of Business
Philadelphia*

SUMMARY

by Peter Williamson

*Amos Tuck School of Business Administration
Dartmouth College, Hanover, New Hampshire*

Jeremy Siegel began his presentation with a table of U.S. market historical returns and excess equity returns for five time periods. **Table 1** provides returns for two very long periods, from the 1800s to September 30, 2001, for three subperiods making up the long periods, and for five post-World War II periods. What is most noteworthy in Table 1 is the geometric (compounded) average real return on stocks of close to 7 percent for the long periods, for both of the major subperiods, and for the 1946–2001 period. Equally significant are the wide deviations above and below 7 percent over quite long periods

after World War II, especially since 1982. The geometric average for 1982–1999 was 13.6 percent, and Siegel concluded that this high average return has influenced the high expectations of today's investors, many of whom have little experience of the pre-1982 period.

Table 1 indicates that average real U.S. T-bond returns fell over the years until the post-1982 period, when very high returns resulted from a decline in interest rates. The 1926–2001 period produced a 2.2 percent average real bond return, biased downward by unexpected inflation in the 1960s and 1970s. Siegel observed that TIPS were priced originally in 1997 at about 3.375 percent, with the yield later rising to about 4 percent, and are now down to about 3 percent.¹ This pricing is close to the 200-year average real return on bonds.

¹ TIPS are Treasury Inflation-Protected Securities; these securities are now called Treasury Inflation-Indexed Securities.

Table 1. Historical Returns and Equity Premiums, 1802–September 2001

Period	Real Return						Stock Excess Return over			
	Stocks		Bonds		Bills		Bonds		Bills	
	Comp.	Arith.	Comp.	Arith.	Comp.	Arith.	Comp.	Arith.	Comp.	Arith.
1802–2001	6.8 %	8.4 %	3.5 %	3.9 %	2.9 %	3.1 %	3.4 %	4.5 %	3.9 %	5.3 %
1871–2001	6.8	8.5	2.8	3.2	1.7	1.8	3.9	5.3	5.0	6.6
<i>Major subperiods</i>										
1802–1870	7.0 %	8.3 %	4.8 %	5.1 %	5.1 %	5.4 %	2.2 %	3.2 %	1.9 %	2.9 %
1871–1925	6.6	7.9	3.7	3.9	3.2	3.3	2.9	4.0	3.5	4.7
1926–2001	6.9	8.9	2.2	2.7	0.7	0.8	4.7	6.2	6.1	8.0
<i>Post World War II</i>										
1946–2001	7.0 %	8.5 %	1.3 %	1.9 %	0.6 %	0.7 %	5.7 %	6.6 %	6.4 %	7.8 %
1946–1965	10.0	11.4	–1.2	–1.0	–0.8	–0.7	11.2	12.3	10.9	12.1
1966–1981	–0.4	1.4	–4.2	–3.9	–0.2	–0.1	3.8	5.2	–0.2	1.5
1982–1999	13.6	14.3	8.4	9.3	2.9	2.9	5.2	5.0	10.7	11.4
1982–2001	10.2	11.2	8.5	9.4	2.8	2.8	1.7	1.9	7.4	8.4

Note: Comp. = compound; Arith. = arithmetic.

Sources: Data for 1802–1871 are from Schwert (1990); data for 1871–1925 are from Cowles (1938); data for 1926–2001 are from the CRSP capitalization-weighted indexes of all NYSE, Amex, and Nasdaq stocks. Data through 2001 can be found in Siegel (2002).

Real returns on T-bills averaged 2.8 percent from 1982 to September 30, 2001—a surprisingly high return for those who were accustomed to the popular position a few years ago that bills offered a zero real rate.

The equity excess return, over both bonds and bills, from 1982 to 1999 and from 1926 to 2001 was much higher than it had been for the long periods, and Siegel commented that the 3–4 percent range that characterized the longer periods was probably reasonable for the long term.

Figure 1 shows the historical P/E of the equity market (calculated from the current price and the last 12 months of reported earnings) for 1871 through September 2001. The collapse of earnings recently pushed the ratio up to 37, past the high of 1999. The average P/E over 130 years was only 14.5. Siegel noted that huge losses in only a few technology companies accounted for a lot of this valuation change. Real stock returns have been biased upward with the rise in P/Es. If the market's P/E were to return to the historical (since 1871) average overnight, the real return on equities would fall 50 bps. And if the P/E had always remained at its average level, without reinvestment of the dividends that actually were paid, real returns would be 115 bps lower than where they are today.

Siegel offered three reasons for rising P/E multiples. First is declining transaction costs, which could

have accounted for 2 percent a year in the 19th and early 20th centuries and are presently perhaps as low as 0.2 percent for a one-way trade. Second is declining real economic risk. And third is investors learning more about the long-term risk characteristics of common stocks, especially investors realizing that there are periods of significant undervaluation.

Table 2 shows the relationships among real stock returns, P/Es, earnings growth, and dividend yields. For 130 years, the real stock return, averaging 7 percent, has been almost exactly the earnings yield (reciprocal of the P/E). The periods before and after World War II show close to the same 7 percent. Faster post-WWII earnings growth matches the decline in the dividend yield and the rise in retained earnings. Siegel noted that this long-term relationship between payout and growth is in accord with theory, but over short periods, the change in earnings growth does not always accompany a change in dividend yield.


The link between P/E and real returns is given by

$$\text{Expected future real returns} = \frac{E}{P} + g \left(1 - \frac{RC}{MV} \right),$$

where

- E/P = earnings yield, the inverse of the P/E
- g = real growth
- RC = replacement cost of capital
- MV = market value of capital
- RC/MV = book-to-market value, or 1/Tobin's q

Figure 1. Historical Market P/E, 1871–2001



Note: Ending month for 2001 is September.

HISTORICAL RESULTS I

Table 2. Historical Growth Rates, 1871–September 2001

Period	Real Stock Return	Average P/E	Inverse of Average P/E	Real Earnings Growth	Dividend Yield	Real Dividend Growth	Real Capital Gains	Average Payout Ratio
1871–2001	7.06 %	14.45	6.92 %	1.27 %	4.66 %	1.09 %	2.17 %	62.24 %
1871–1945	6.81	13.83	7.23	0.66	5.31	0.74	1.32	70.81
1946–2001	7.38	15.30	6.54	2.08	3.78	1.57	3.32	50.75

Tobin's q is currently about 1.2, and the long-run growth rate, g , is about 3 percent, so the term $g[1 - (RC/MV)]$ adds about 0.5 percentage point to the E/P term. At a P/E of 20, appropriate for today, the

expected real return is about 5.5 percent. At a P/E of 25, it is 4.5 percent. With the TIPS return at about 3 percent and a P/E of 20 to 25, Siegel's equity risk premium is about 2 percent (200 bps).

Historical Results II

Bradford Cornell

*University of California
Los Angeles*

The basic investment and constant-growth models, used with some justifiable simplifying assumptions about the U.S. market, indicate that the earnings growth rate cannot be greater than the GNP growth rate because of political forces and that the expected return, or cost of capital, in the long run should unconditionally be about 1.5 times the dividend-to-price ratio plus GNP growth. Adding reasonable assumptions about inflation produces a finding that equity risk premiums cannot be more than 3 percent (300 bps) because earnings growth is constrained by the real growth rate of the economy, which has been in the 1.5–3.0 percent range. In a consideration of today's market valuation, three reasons for the high market valuations seem possible: (1) stocks are simply seen as less risky, (2) valuation of equities is fundamentally determined by taxation, or (3) equity prices today are simply a mistake. A research question that remains and is of primary interest is the relationship between aggregate stock market earnings and GNP.

The very basic investment and constant-growth models from introductory finance courses can be used to interpret the long-run unconditional historical data on returns. So, let's begin with the basic model:

$$\frac{E_{t+1}}{E_t} = 1 + [(b)(ROE)],$$

where

E = earnings

b = the retention rate

ROE = return on equity

So that, with investment at time t denoted by I_t ,

$$ROE = \frac{E_{t+1} - E_t}{I_t}$$

and

$$b = \frac{I_t}{E_t};$$

therefore, the growth rate of earnings is

$$(b)(ROE) = \frac{E_{t+1} - E_t}{E_t}.$$

This model implies that the growth rate in earnings is the retention rate times the return on equity, $(b)(ROE)$. In discussing the models, I would like to stress an important point: If you are interpreting the growth in earnings as being the retention rate times the return on equity, you have to be very careful when you are working with historical data. For example, does the retention rate apply only to dividends or to dividends and other payouts, such as share repurchases? The distinction is important because those proportions change in the more recent period. And if you make that distinction, you have to make a distinction between aggregate dividends and per share dividends because the per share numbers and the aggregate numbers will diverge. In working with the historical data, I have attempted to correct for that aspect.

What simplifying assumptions can be made to work with the unconditional data? I have made some relatively innocuous simplifying assumptions. First, that b should adjust until the cost of capital equals the ROE at the margin. To be very conservative, therefore, I will assume that the ROE equals the cost of capital, or expected returns, in the aggregate. The problem that arises is: What if the retention rate times the cost of capital (that is, the minimal expected return on equity), bk , is greater than GNP growth? The second assumption deals with this possibility: I assume bk cannot be greater than GNP growth because political forces will come into play that will limit the ROE if earnings start to rise as a fraction of GNP.

The relationship between aggregate earnings and GNP is one of the research questions that I have been unable to find interesting papers on—perhaps because I have not searched well enough—but I want to bring up the subject to this group. It seems to me that if aggregate earnings start to rise, and Robert Shiller mentioned several reasons why it can happen [see the “Current Estimates and Prospects for Change” session], then tax rates can change, antitrust regulation can change (one of Microsoft’s problems probably was that it was making a great deal of money, which is an indication that some type of regulation may be necessary), labor regulation can change, and so forth. And these variables can change *ex post* as well as *ex ante*. So, once a company starts making superior returns using a particular technology, the government may step in *ex post* and limit those returns. The critical research question is how earnings relate to GNP.

The constant-growth model is

$$P = \frac{D}{k - g}$$

or

$$k = \frac{D}{P} + g,$$

where

P = price

D = dividends

k = cost of capital

g = growth rate

What I am going to do is just an approximation because I am going to work with aggregate, not per share, data. I am going to assume that total payouts are 1.5 times dividends.¹ Payouts will probably be lower in the future, but if I work with aggregate

¹This choice is based on recent findings by Jagannathan, Stephens, and Weisbach (2000) that we are seeing significant payouts today.

payouts, then g should be the growth rate in aggregate potential payouts, which I will characterize as earnings.

One of the implications of the simplifying assumptions I have made, and it relates to the data that Jeremy Siegel just produced [“Historical Results I”], is that the expected returns on stocks should be equal to the earnings-to-price ratio. (In the more complicated equations, you have situations in which the ROE is not exactly equal to expected returns, but for my long-run data, the simplifying assumption that earnings yield equals the expected ROE is fine.) So, with these assumptions,

$$\begin{aligned} P &= \frac{D}{k - g} \\ &= \frac{D}{k - bk} \\ &= 1 - (b) \left(\frac{E}{1 - b} \right) (k) \\ &= \frac{E}{k} \end{aligned}$$

or

$$k = \frac{E}{P}.$$

A further implication is that if g is constrained to be close to the growth of GNP, then it is reasonable to substitute GNP growth for g in the constant-growth model. The implication of this conclusion is that the expected return, or cost of capital, in the long run should unconditionally be about 1.5 times the dividend-to-price ratio plus GNP growth:

$$k = 1.5 \frac{D}{P} + \text{GNP growth}.$$

With this background, we can now look at some of the data.

Earnings and GNP

Figure 1 allows a comparison of dividends/GNP and (after-tax) earnings/GNP for 1950 through July 2001.² The data begin in 1950 because Fama believed that the data before then were unreliable. Figure 1 shows that, historically, earnings have declined as a fraction of GNP in this period. My assumption that earnings keep up with GNP works from about 1970 on, but I am looking at the picture in Figure 1 in order to make that conclusion. The ratio of earnings to GNP depends on a lot of things: the productivity of labor, capital, the labor-to-capital ratio, taxes, and (as I said earlier) a host of political forces. Figure 1 shows that earnings have, at best, kept up with GNP.

²These data were provided by Eugene Fama, who attributed them to Robert Shiller.

Figure 1. S&P 500 Earnings and Dividends to GNP, 1950–July 2001




Table 1 gives the arithmetic average data for growth rates in GNP, earnings, and dividends for two periods: 1951–2000 and 1972–2000. (I used the 1972–2000 period because it mirrors the same period shown in Figure 1.) The earnings growth rates are so much more volatile than the dividend growth rates. And because of the volatility effect on arithmetic averages, GNP and earnings exhibit very similar growth rates from the early 1970s to the present. Dividends (and Table 1 shows the growth rate of actual dividends, not payouts) have grown much less than earnings for two reasons: First, dividends are less volatile, and second, dividend substitution is occurring. Corporations are not providing shareholders the same constant fraction of earnings (in the form of dividends) that they were in the past.

Despite the 1972–2000 data, it seems to me that earnings are not going to grow as fast as or faster than GNP in the future. This notion seems to be consistent with long-term historical data, and it fits my view of how politics works on the economy. If you accept that notion, it has immediate implications for the future.

Table 1. Historical Growth Rates of GNP, Earnings, and Dividends: Two Modern Periods

Period/Measure	GNP	Earnings	Dividends
<i>1951–2000</i>			
Mean	3.21 %	2.85 %	1.07 %
Standard deviation	2.89	14.29	4.13
<i>1972–2000</i>			
Mean	2.62 %	3.79 %	0.96 %
Standard deviation	2.94	15.72	3.58

Note: Growth rates for earnings and dividends are based on aggregate data.

First, under any reasonable underlying assumptions about inflation, equity risk premiums cannot be much more than 3 percent (300 bps) because the earnings growth rate is constrained unconditionally in the long run by the real growth rate of the economy, which has been in the range of 1.5–3.0 percent. Second, as **Table 2** shows, for an S&P level of about 1,000, you simply cannot have an equity risk premium any higher than 2 percent, 2.5 percent, or (at most) 3 percent.

Table 2. Value of the S&P 500 Index Given Various Real (Earnings or GNP) Growth Rates and Equity Risk Premiums

Real Growth Rate	Equity Risk Premium						
	2.0 %	2.5 %	3.0 %	4.0 %	5.0 %	6.0 %	7.0 %
1.5 %	845	724	634	507	423	362	317
2.0 %	1,014	845	724	563	461	390	338
2.5 %	1,268	1,014	845	634	507	423	362
3.0 %	1,690	1,268	1,014	724	563	461	390

Assumptions: Inflation = 3 percent; long-term risk-free rate = 5.5 percent; payout = 1.5(S&P 500 dividend). The S&P 500 dividend used in the calculation was \$16.90, so $P = 1.5(\$16.90)/(k - g)$, where $k = 5.5$ percent (the risk-free rate minus 3 percent inflation plus the risk premium) and g = real growth rate.

Valuation

Why is the market so high? As an aside, and this concern is not directed toward our topic today of the equity risk premium, but I think it is an interesting question: Why is the market where it is today relative to where it was on September 10 or September 9 or just before the events of September 11, 2001? The market then and now is at about the same level. Almost every economist and analyst has said that the September 11 attacks accelerated a recession, that they changed perceptions of risk, and so forth. It is curious to me that such a situation does not seem to be reflected in market prices.

But in general, why is the market so high? I believe three possible explanations exist. One idea, and I consider it a “rational” theory, is that stocks are simply seen as less risky than in the past. I do not know whether the behavioral theories are rational or not, in the sense that prices are high because of behavioral phenomena that are real and are going to persist. If so, then those phenomena—as identified by Jeremy Siegel and Richard Thaler [see the “Theoretical Foundations” session]—are also rational. In that case, the market is not “too high”; it is not, in a sense, a mistake. It is simply reflecting characteristics of human beings that are not fully explained by economic theories.

Another rational explanation has been given less attention but is the subject of a recent paper by McGrattan and Prescott (2001). It is that the valuation of equities is fundamentally determined by taxation. McGrattan and Prescott argue that the move

toward holding equities in nontaxable accounts has led to a drop in the relative tax rate on dividends. Therefore, stock prices should rise relative to the valuation of the underlying capital and expected returns should fall. This effect is a rational tax effect.

Both this theory and the theory that stocks are now seen as less risky say that the market is high because it should be high and that, looking ahead, equities are going to have low expected returns, or low risk premiums—about 2 percent—but that investors have nothing to worry about.

The final explanation, which I attribute to John Campbell and Robert Shiller, focuses on the view that equity prices today are simply a mistake. (I suppose mistakes are a behavioral phenomenon, but presumably, they are not as persistent as an underlying psychological condition.) Now, when people realize they have made a mistake, they attempt to correct the behavior. And those corrections imply a period of *negative* returns from the U.S. equity market before the risk premium can return to a more normal level.

Closing

To close, I want to repeat that, to me, the fundamental historical piece of data that needs more explanation is the relationship between the aggregate behavior of earnings and GNP—what it has been in the past and what it can reasonably be going forward. This relationship is interesting, and I look forward to hearing what all of you have to say about it. In my view, it is the key to unlocking the mystery of the equity risk premium’s behavior.

Historical Results II

Bradford Cornell

University of California
Los Angeles

SUMMARY

by Peter Williamson

Amos Tuck School of Business Administration
Dartmouth College, Hanover, New Hampshire

To interpret long-run unconditional features of historical returns, Bradford Cornell began with the following basic model:

Earnings growth = $(b)(ROE)$,

where b is the rate at which earnings are retained and ROE is return earned on equity. He noted that we have to be careful when working with historical data in this model. For example, does payout apply only to dividends or to dividends and other payouts, such as share repurchases? And we need to distinguish between aggregate dividends and per share dividends. The two have been diverging.

Now, b should adjust until ROE at the margin equals k , the cost of capital. Cornell assumed that $k = ROE$ in the aggregate, but a critical question is how earnings relate to GNP (see **Figure 1**). What if

bk is greater than GNP growth? Cornell assumed that political forces—such as taxation, antitrust laws, and labor regulations—would affect *ex ante* and *ex post* returns in such a way as to bring about

$$(b)(ROE) = bk \leq \text{GNP growth.}$$

The constant-growth model is

$$P = \frac{D}{k - g}$$

or

$$k = \frac{D}{P} + g$$

where

P = price

D = dividends

k = cost of capital


g = growth rate

Because D is equal to $E(1 - b)$ and g is equal to bk , the constant-growth model becomes, in real terms,

$$P = \frac{E}{k}$$

or

Figure 1. S&P 500 Earnings and Dividends to GNP, 1950–July 2001



$$k = \frac{E}{P}$$

Cornell had so far been working with aggregates, but share repurchases and other nondividend cash flows between companies and their shareholders should be considered. So, he assumed that the total of cash distributions is approximately $1.5D$.

Finally, if g is constrained to be close to GNP growth, then $k = 1.5(D/P) + \text{GNP growth}$.

Table 1 shows that since 1950, aggregate S&P 500 Index earnings and dividends have both grown less than GNP, although from 1972 to 2000, earnings actually grew faster. (Earnings may appear to have kept up with or even exceeded GNP because of the high volatility of the earnings, which leads to high arithmetic average rates of growth for the same geometric averages.) The dividend growth rates have been lower because of falling payout ratios. The picture conveyed to Cornell is that earnings growth will not exceed GNP growth in the future. (The relationship of earnings to GNP is an interesting measure

Table 1. Historical Growth Rates of GNP, Earnings, and Dividends: Two Modern Periods

Period/Measure	GNP	Earnings	Dividends
<i>1951–2000</i>			
Mean	3.21 %	2.85 %	1.07 %
Standard deviation	2.89	14.29	4.13
<i>1972–2000</i>			
Mean	2.62 %	3.79 %	0.96 %
Standard deviation	2.94	15.72	3.58

Note: Growth rates for earnings and dividends are based on aggregate data.

having to do with, among other things, the productivity of labor and capital.)

Finally, putting together an inflation assumption of 3 percent, a long-term nominal risk-free rate of 5.5 percent, and the relationships developed previously produces **Table 2**. An example of the calculations for Table 2 under the assumptions given in the table is as follows: At real growth of 3 percent and with a risk premium of 2.5 percent, $P = [1.5(\$16.90)] / (0.055 - 0.03 + 0.025 - 0.03) = \$1,268$. What Table 2 indicates is that as long as g is limited by GNP growth of 1.5–3.0 percent, the equity risk premium must be no more than about 3 percent to be consistent with an S&P 500 of about 1,000.

Cornell asked why, in general, is the market so high? (In particular, he questioned why the market is currently at the level of pre-September 11, 2001, if, as so many say, the events of that date accelerated a recession and changed perceptions of risk.) One explanation is that investors see the market generally as less risky than in the past. Cornell found that explanation rational. Another rational explanation is that the value of equities is fundamentally determined by taxation. Perhaps the market's level is explained by human behavior that is rational but for which we have no explanation. Both propositions imply that there is nothing wrong with current prices. Still, another explanation is that equity prices are a mistake and that a downward correction will produce negative returns before a normal risk premium prevails.

A key subject on which we might focus is the relationships among aggregate earnings, GNP, and other economic variables.

Table 2. Value of the S&P 500 Index Given Various Real (Earnings or GNP) Growth Rates and Equity Risk Premiums

Real Growth Rate	Equity Risk Premium						
	2.0 %	2.5 %	3.0 %	4.0 %	5.0 %	6.0 %	7.0 %
1.5 %	845	724	634	507	423	362	317
2.0 %	1,014	845	724	563	461	390	338
2.5 %	1,268	1,014	845	634	507	423	362
3.0 %	1,690	1,268	1,014	724	563	461	390

Assumptions: Inflation = 3 percent; long-term risk-free rate = 5.5 percent; payout = $1.5(\text{S\&P 500 dividend})$. The S&P 500 dividend used in the calculation was \$16.90, so $P = 1.5(\$16.90) / (k - g)$, where $k = 5.5$ percent (the risk-free rate minus 3 percent inflation plus the risk premium) and $g = \text{real growth rate}$.

Historical Results: Discussion

Ravi Bansal (Moderator)

Robert Arnott

Clifford Asness

John Campbell

Peng Chen, CFA

Bradford Cornell

William Goetzmann

Campbell Harvey

Roger Ibbotson

Martin Leibowitz

Rajnish Mehra

Thomas Philips

William Reichenstein, CFA

Stephen Ross

Robert Shiller

Jeremy Siegel

Kevin Terhaar, CFA

Richard Thaler

RAVI BANSAL (Moderator)

I would like to make a couple of observations. One aspect that we could consider is the time-series evidence on aggregate consumption volatility. I am thinking of consumption as a way to measure economic uncertainty in the data, but it can be done by other means as well. The time-series evidence suggests that a decline in conditional volatility has without doubt occurred over the past 40 years or so. This reduced volatility suggests that there should be some decline in risk premiums. Another aspect that could be considered, which Steve Ross mentioned earlier, is that much of the risk premium discussion draws on the cross-sectional evidence. It is where a lot of the bodies are buried in terms of understanding where risks are coming from.

We heard some debate in the first session ["Theoretical Foundations"] about whether consumption models are plausible or not, and my view is that consumption data are not in a usable form for explaining the cross-sectional differences, although there may be new evidence in this regard. The consumption models can actually go a long way, however, in explaining the difference in the risk

premiums on different assets. In fact, in "Consumption, Dividends, and the Cross-Section of Equity Returns" (Bansal, Dittmar, and Lundblad 2001), we show that if you take the earnings growth or the dividend growth of different portfolios and regress actual growth on historical (say, the past 25–30 years) consumption growth smoothed for 12 or 14 quarters, and if you consider (what has almost become the industry benchmark) 10 portfolios composed on the basis of size, 10 on momentum, and 10 on the book-to-market ratio, you will see that the regression coefficient almost entirely lines up with the *ex post* excess returns on these different assets. So, for example, the regression coefficient of extreme "loser" momentum portfolios is negative and that of "winner" portfolios is strongly positive. The value stocks have a very high exposure to the consumption growth rate, and what I call the loser value stocks—that is, the growth stocks—have a low exposure, which maps the differences in equity premiums also. So, there is a link between consumption and risk premiums, which creates a *prima facie* case for aggregate economic uncertainty, defined as consumption, being a very useful measure.

The cross-sectional evidence also highlights that what determines the risk premium on an asset is "low-frequency" movements (long-run growth prospects) and the exposure of different portfolios to them. Long-run growth prospects are the key source of risk in the economy.

Still, a puzzle remains because the equity market risk premiums have decreased—to 2 percent, 2.5 percent, or so on—and of course, people disagree about what the risk premium is. It seems to me that the right way to approach the equity risk premium puzzle is through the Sharpe ratio on the market. If we argue that the risk premium has fallen, then the Sharpe ratio is quite likely to have fallen also.

CLIFFORD ASNESS: If I understood correctly, Jeremy Siegel was saying that Rob Arnott and I were picking up a short-term mean-reversion effect that is not relevant over the long term. I would like to make two points: First, we were forecasting over several decades and found a pretty strong negative relationship between the retention rate and real earnings growth. So, Jeremy, if this relationship reverses itself in the longer term, we should find a very, very strong positive relationship later. Yes? Second, in the draft of our paper (Arnott and Asness 2002), which has

only been seen by Rob, me, and a few people we trusted not to laugh at us, we tested the relationship against other proxies for pure, univariate mean reversion in earnings growth—prior growth, growth versus a 20-year average—added to the equation. We still found over a 10-year horizon (we would like to have used a longer horizon but were trying to avoid having too few periods) that the relationship is very negative. Therefore, I have a hard time believing that over longer periods the relationship is going to be very positive. We did find that simple measures of mean reversion and earnings do not knock out the relationship. I am curious about the data you were using and what you are citing in the longer term. Maybe we can reconcile the apparent differences.

JEREMY SIEGEL: Well, I did not run the tests that you did. I just know that there is very strong evidence from cycles. In recessions, the payout ratio goes very high because companies choose to maintain the same level of dividends they were paying before the recession, and earnings drop. Then, subsequent growth in real earnings is very high because it is happening relative to the slow or negative growth experienced during the recession. The same phenomenon, but in the opposite direction, occurs during and after an economic boom. For these reasons, I found in the two long periods, 1871–1945 and then 1946–2000, that the decrease in the dividend yield during each period was matched by an increase in real earnings growth [see Siegel’s Table 2]. The result is the same approximate 7 percent real return in the later period as in the earlier period, which is comforting from a theoretical point of view. Otherwise, we would have to turn to such theories as that “companies that retain more earnings must be totally wasting them because the companies do worse after the earnings retention.” That theory is very much a concern.

JOHN CAMPBELL: I want to focus attention on an issue that is in Jeremy Siegel’s tables but which he didn’t talk about in his presentation—the geometric versus the arithmetic average. This issue is one that causes people’s eyes to glaze over. It seems a pedantic thing, like worrying about split infinitives—the sort of thing that pedantic professors do but other people shouldn’t bother about. But it is actually an important issue for risky assets because the difference between the arithmetic and the geometric average is on the order of about half the variance, which for stocks, is about 1.5–2.0 percent. That’s a big difference, and it shows up in Jeremy’s tables very clearly. So, when we’re bandying about estimates of the equity premium and we say, “Maybe it’s 2 percent; maybe it’s 3

percent,” clearly the difference between these two averages is large relative to those estimates.

Which is the right concept, arithmetic or geometric? Well, if you believe that the world is identically and independently distributed and that returns are drawn from the same distribution every period, the theoretically correct answer is that you should use the arithmetic average. Even if you’re interested in a long-term forecast, take the arithmetic average and compound it over the appropriate horizon. However, if you think the world isn’t i.i.d., the arithmetic average may not be the right answer.

As an illustration, think about a two-lane highway to an airport. Suppose that to increase traffic capacity, you repaint the highway so that it has three, narrower lanes. Traffic capacity is thus increased by 50 percent. But suppose the lanes are now too narrow, causing many accidents, so you repaint the highway with only two lanes. Arithmetically, the end result appears to be a great success because the net effect is an increase in capacity. A 50 percent increase in capacity has been followed by only a 33.3 percent decrease. The arithmetic average of the changes is +8.5 percent. So, even though you’re back to your starting point, you delivered, on average, an 8.5 percent increase in traffic capacity. Obviously, that’s absurd. In this case, the geometric average is the right measure. The geometric average calculates a change in capacity to be zero, which is the correct answer; nothing has been accomplished with the lane rearrangement and reversal.

The difference between the i.i.d. case and the highway story is that in the highway story, you have extreme negative serial correlation. You could get to –33.3 percent in the end only by having had the +50 percent and –33.3 percent occur on a higher base than +50 percent. So, the geometric average is the correct measure to use in an extreme situation like the highway illustration.

I think the world has some mean reversion. It isn’t as extreme as in the highway example, but whenever any mean reversion is observed, using the arithmetic average makes you too optimistic. Thus, a measure somewhere between the geometric and the arithmetic averages would be the appropriate measure.

BRADFORD CORNELL: You see that difference in the GNP and earnings data. Although the ratio of earnings to GNP is falling from 1972 on [see Cornell’s Table 1], the growth rate of earnings is higher as an *arithmetic* mean precisely for the reason you suggest.

CAMPBELL: Right, right. Mean reversion has the effect of lowering the variance over long horizons, which is, of course, a major theme of Jeremy Siegel’s

work. And you could imagine taking the geometric average and then adding half of long-term variance to get an appropriate long-term average.

SIEGEL: That's a good point. You discussed in your new book with Lewis Viceira (Campbell and Viceira 2002) whether we should use the arithmetic or the geometric average and that when mean reversion occurs, we perhaps have more reason to use the geometric average. I've found in my data that at 30-year horizons, the standard deviation is about half the number that i.i.d., random walk theory would predict. So, you can actually add half the variance to the geometric average and use that number as the appropriate arithmetic risk premium on long horizons.

CAMPBELL: It was striking that you did focus your presentation on the geometric average. A lot of the other calculations that have been presented here today evolve out of these deterministic models in which no distinction is made between geometric and arithmetic calculations. But I think that when you face randomness, as we do in the world, you have to think about this issue.

ROBERT ARNOTT: I had just a quick follow-up to Cliff Asness's question about the link between payout ratios and earnings growth. I think one possible source of the difference that we're seeing is not the time horizon but that, in Jeremy Siegel's work, if I understand correctly, he is looking at the *concurrent* payout ratio versus earnings growth. Cliff Asness and I are looking at *leading* payout ratio versus *subsequent* earnings growth; in effect, we're using the payout ratio as a predictor of earnings growth.

ASNESS: I'll add one thing to that: What Jeremy Siegel is saying is that a high and falling dividend yield is replaced by increased earnings growth over that period. What Rob Arnott and I are saying is that perhaps there is mean reversion but if you look at the start of that period, the high dividend yield was leading to a high payout ratio, which tended to forecast the declining actual earnings growth. So, I think we're actually saying the same thing. That's a limb I'm going to go out on.

CAMPBELL HARVEY: One thing that completely baffles me is the TIPS yield right now. The breakeven inflation rate for 10 years is about 1.2 percent. Brad Cornell showed that valuation table [Cornell's Table 2] with a reasonable assumption of inflation at 3 percent. And Jeremy Siegel's Table 1 showed the historical data in terms of real bond return, which was significantly higher on average than 1.5 percent. It just seems there's something going on with TIPS

that I don't understand. For me, an inflation rate of 1.2 percent over 10 years doesn't seem reasonable.

PENG CHEN: It depends on how you define the equity risk premium. Some define the equity risk premium in relation to the real return earned on TIPS. It's a good observation, but TIPS is a new asset class, started just several years ago. The TIPS market is still immature; the market size is relatively small. So, I'm not sure how much inference you should draw by just looking at the current yield. A current yield of 3 percent doesn't mean that the real interest rate is 3 percent. If you had followed the TIPS market for a while, you probably would have heard rumors that the U.S. Treasury Department is going to suspend issuing TIPS—which would have a huge impact on how TIPS behave in the marketplace. So, we need to be careful when using TIPS as part of the benchmark in trying to calculate the actual risk premium.

SIEGEL: On that issue, I think there is a liquidity issue with TIPS, but it's not that great. I think there's \$70, \$80, \$90 billion worth of TIPS in the market. You can do a trade of fairly decent size at narrow bid-ask spreads. My opinion of what's going on right now is that nominal bonds are seen as a hedge. I think there is fear of deflation in the market. And as in 1929, 1930, and 1931, investors were thinking that if the world markets, such as Japan, were going to be in a bad state, in a deflationary sense, holding nominal assets was the thing to do. So, as a result, the demand for nominal bonds is rising as a hedge against deflation, which will be bad for the economy and for real assets. The difference between TIPS and nominal bonds doesn't measure unbiased expected inflation; there's a negative risk premium in the picture. It is not what we think of as "there's inflation risk so nominal bonds should sell at a higher-than-expected return." I think right now the premium is a negative risk premium as investors use nominal bonds as a hedge against deflationary circumstances in the economy.

STEPHEN ROSS: In all of these computations of the equity risk premium on the stock market, does anyone take into account the leverage inherent in the stock market and the volatility premium that you would get from it? I don't have a clue about the empirical size of that premium. Can someone help me?

MARTIN LEIBOWITZ: I can. If you take the formulas that have been discussed today and translate them to assume a particular risk premium on unlevered assets, you can see how that premium translates into the typical level of leverage in the equity markets. You find that it is exactly what you'd expect. The risk premium that you actually see in the market reflects

the leverage that is endemic in the equity market, and if you back out that premium to find the risk premium on unlevered assets, you find that the premium on unlevered assets is less.

RAJNISH MEHRA: The Sharpe ratio won't change. It's invariant to leverage.

LEIBOWITZ: It's exactly linear.

ROBERT SHILLER: Let's remember correctly the McGrattan and Prescott article (2001) that Brad Cornell mentioned. They use a representative agent model, and they compare the late 1950s and early 1960s with a recent year. And they say that because of 401(k)s and similar vehicles, the tax rate on dividends for a representative agent has fallen—from 50 percent in 1950–1962 to 9 percent in 1987–1999. That fall seems to me like an awfully big drop, and I question whether there could have been such a big drop for the representative investor. I wonder if anyone here has looked carefully at their model? Are they right?

SIEGEL: They use the average investor; they don't use the marginal investor. They say that X percent of assets are in a 401(k), and they equate that amount with the marginal rate. My major criticism of the McGrattan–Prescott paper is that we don't know whether the marginal investor is a taxable investor, which would change their results dramatically.

CORNELL: That criticism doesn't mean their results are wrong. We simply don't know.

SIEGEL: We don't know. But I have a feeling that the marginal investor has a much higher tax rate than the marginal investor used to have.

ROSS: Yes, James Poterba told me that his calculations indicate that 401(k)s have far less tax advantage at the margin than one might think. Because of the tax rate “upon withdrawals,” those vehicles can be dramatically attacked from a tax perspective. If you make a simple presumption that 401(k)s are simply a way of avoiding taxes, you're missing the point.

THOMAS PHILIPS: I'd like to go back to the equation for expected future real returns that Jeremy Siegel attributes to me: Expected future real returns = Earnings yield + $g \times [1 - (\text{Book value}/\text{Market value})]$. It really is an expression for the expected future *nominal* return. When I derived that equation, I derived it in *nominal* terms. In particular, the growth term, g , is nominal, not real, growth (Philips 1999). When you subtract inflation, you have Expected future real returns = Earnings yield + Nominal growth $\times [(1 - \text{Book value}/\text{Market value}) - \text{Inflation}]$; the last two

terms go to approximately zero. You're left with the earnings yield being approximately the real expected return.

In the special case that Brad Cornell talked about, in which the cost of capital and the return on capital are the same, the second term disappears because the book-to-market ratio becomes 1. In that case, the earnings yield is actually the *nominal* expected return. The truth, in practice, lies somewhere in between the two results because some of these quantities will vary with inflation, real interest rates, and the economywide degree of leverage.

The approximation that Brad used is biased up or down depending on where inflation, growth, and the cost of capital relative to the return on capital lie. It's a great first-order approximation, a great historical approximation, but you can be talking about the nominal rate of return instead of the real rate of return when the cost of capital starts coming very close to the return on capital.

SIEGEL: Well, I disagree with you. In your slides, the earnings yield—if you're in equilibrium and book value equals market value equals replacement cost—is an estimate of the real return, not the nominal return. Your equation is extraordinarily useful, but I think we do have to interpret it as the real return.

ROGER IBBOTSON: I'd like to say something about Brad Cornell using aggregate calculations to get an estimate of the equity risk premium. I did some work on aggregate calculations in a paper I wrote with Jeffrey Diermeier and Laurence Siegel in 1984. Relating to merger and acquisition activities, we looked at how best to use cash: For example, do you use cash for dividends or share repurchases? (You could take the same approach for investing in projects.) When you look at which data to use in the context of cash mergers or acquisitions, you can see that the per share estimates are going to be very different from the aggregate estimates because you're buying other companies on a per share basis. Thus, EPS can grow much faster than aggregate corporate earnings.

CORNELL: That's why I like looking at aggregate earnings; it's the whole pot, and you're not as concerned about how things are moving around within the pot or being paid out to shareholders. But even looking at aggregate earnings, and this is based on Bob Shiller's data series going back to 1872, the earnings don't keep up with GNP, despite the greater volatility of earnings; even the arithmetic averages are less. Can you explain that phenomenon? What does it imply for the future?

SHILLER: The national income and product account (NIPA) earnings keep up a lot better. So, it's probably because earnings in the market indexes are not representing the new companies that come into the economy and existing companies' earnings are growing at a slower rate.

SIEGEL: I looked at it very closely. The trend in the ratio of NIPA profits to GDP is virtually zero, the mean being 6.7 percent. You can do a linear regression—any regression—and you get a trend of absolutely zero: The ratio of NIPA profits to GDP has remained constant. Aggregate S&P 500 Index profits have slipped because the S&P 500 back in the 1950s and 1960s represented a much higher percentage of the market's value than it does today. You can look at both aggregate S&P 500 profits and aggregate NIPA profits and see the trends.

MEHRA: I found the same thing in my 1998 paper. The ratio of aggregate cash flows to national income (NI) is essentially trendless. In the afternoon, I'll be talking about the difference when you look at stock market valuation relative to national income [see the "Current Estimates and Prospects for Change" session]. That ratio fluctuates from about $2 \times \text{NI}$ to about $0.5 \times \text{NI}$, whereas cash flows, which are the input for all these valuation models, are trendless relative to NI.

KEVIN TERHAAR: I want to go back to the representative investor or the marginal investor and Brad Cornell's first "rational" reason that the market might be high—that stocks are seen as less risky. One thing that hasn't been brought up is that all the discussions so far have focused primarily on the U.S. equity market. To the extent that the marginal investor looks at U.S. equities in the context of a broader portfolio (as opposed to looking at them only in a segmented market), the price of risk (or the aggregate Sharpe ratio) can stay the same while the equity premium for U.S. equities can fall. As the behavior of investors becomes less segmented—as they become less apt to view assets in a narrow or isolated manner—the riskiness of the assets can decline. Risk becomes systematic rather than total, and as a result, the compensation for risk falls commensurately.

WILLIAM GOETZMANN: I have a related comment in reference to Brad Cornell's presentation. An interesting aspect was his reference to changes in diversification of individual investors. There's not much empirical evidence on this issue, but it's interesting because we did have a boom in mutual funds through the 1980s and 1990s, with investors becoming more diversified. And the result was that the volatility of

their equity portfolios dropped. We saw a similar trend in the 1920s, at least in the United States, through much growth in the investment trusts.¹ We think of trusts as these terrible entities that we clamped down on in the 1930s, but nevertheless, they did provide diversification for individual investors. So, maybe there is some relationship between the average investor's level of diversification and valuation measures of the equity premium.

It's hard to squeeze much more information out of the time-series data because we don't have many booms like I just described. But we might get something from cross-sectional studies—looking internationally—because we have such differences in the potential for investors in each country to diversify—different costs associated with diversification and so forth. So, maybe we could find out something from international cross-sectional data.²

CAMPBELL: On the diversification issue, I have a couple of cautionary notes. First, I think that diversification on the part of individual investors probably is part of this story, but what matters for pricing ought not to be the diversification of investors with investors equally weighted but with investors *value weighted*. Presumably, the wealthy have always been far more diversified than the small investor. So, if small investors succeed in diversifying a bit more, it may not have much effect on the equity premium.

Second, you mentioned the trend toward increased diversification in recent years. There has also been a trend toward increased idiosyncratic risk in recent years. So, although marketwide volatility has not trended up, there has been a very powerful upward trend since the 1960s in the volatility of a typical, randomly selected stock. So, you *need* to be more diversified now in order to have the same level of idiosyncratic risk exposure as before 1960. It's not clear to me whether the increase in diversification of portfolios has outstripped that other trend or merely kept pace with it.

ROSS: It's not at all obvious to me that the wealthy are more diversified. The old results from estate tax data I found are really quite striking. Keep in mind that the data contain survivorship bias and that the rich got wealthy by owning a company that did well, but as I remember, the mean holding of the wealthy is about four stocks, which is really quite small. Conversely, if you look at the less wealthy investor, many of their assets are tied up in pension plans,

¹ Investment trusts existed solely to hold stock in other companies, which frequently held stock in yet other companies.

² For a discussion of long-term equity risk premiums in 16 countries, see Dimson, Marsh, and Staunton (2001).

where the diversification—even in defined-benefit plans—is subtle and not easy to detect. The same can be said for Social Security.

SIEGEL: I think we should also keep in mind the absolutely dramatic reduction in the cost of buying and selling stocks. Bid–ask spreads are sometimes pennies for substantial amounts of stocks, and transaction costs have decreased virtually to zero. I would think that, even with the increase in idiosyncratic risk, if individual investors *want* to diversify (leaving aside the question of whether they want to diversify or pick stocks), they can do so at a much lower cost today than they could, say, 20 or 30 years ago.

BANSAL: So, your argument for the falling equity premium would be that the costs have gone down more for equities than for bonds?

SIEGEL: Yes.

ASNESS: We still see many investors with tremendously undiversified portfolios. There are psychological biases and errors that can lead to a lack of diversification; we haven't had a rush to the Wilshire 5000 Total Market Index.

RICHARD THALER: To follow up, I want to point out that research on the prevalence of ownership of company stock in 401(k) plans indicates that it's quite high—in some companies, shockingly high. At Coca-Cola, for example, at one time, more than 90 percent of the pension assets were in Coca-Cola stock. The same pattern was common in the technology companies. Talk about investments being undiversified *and* positively correlated with human capital! These situations are very risky.

ASNESS: Have you ever tried to convince an endowment started by one family that what they should really do is diversify?

THALER: Right, right.

ASNESS: You never succeed.

THALER: Research on the founders of companies indicates that they hold portfolios with very low returns and very high idiosyncratic risk.

ASNESS: But they had *had* very high returns at some point.

THALER: Right.

PHILIPS: I'd like to re-explore the earnings versus GDP question. Rob Arnott and Peter Bernstein (2002) find that per share earnings grow more slowly

than the economy for a very simple reason: A large chunk of the growth of the economy is derived from new enterprises, and therefore, the growth in earnings per dollar of capital will be inherently lower than the growth of earnings in the entire economy. Their empirical result is that per share earnings grow at roughly the same rate as per capita GDP. Let's call that the rate of growth of productivity. I, on the other hand, am much more comfortable with the notion of EPS growing at roughly the same rate as the economy as a whole. Why? Because the old economy spins off dividends that it cannot reinvest internally. Those dividends, in turn, can be invested in the new economy, which allows you to capture the growth in the new economy. In effect, you have a higher growth rate and a lower dividend yield, and your per share earnings keep growing at roughly the same rate as the economy as a whole. Do you have a take on that, Jeremy? Do you have an instinctive feel for whether we're missing something here or not?

SIEGEL: If companies paid out all their earnings as dividends (with no reinvestment or buying back of shares) and because (based on the long-run-growth literature) the capital output ratio is constant, then EPS would not grow at all. You would have new shares as the economy grew, through technology or population growth, because companies would have to float more shares over time to absorb new capital. But EPS wouldn't really grow at all. What happens, of course, is that the companies withhold some of their earnings for reinvestment or buyback of shares, which pushes EPS upward. If the earnings growth also happens to be the rate of productivity growth or GDP growth, I think it's coincidental, not intrinsic.

IBBOTSON: I have done work on the same subject, and I agree.

WILLIAM REICHENSTEIN: I have a concern. If you're buying back shares, EPS grow (corporate earnings don't necessarily grow, but earnings per share do). The argument that when companies reinvest their earnings rather than paying out their earnings to shareholders they must be wasting some of that money just doesn't jibe with the reality that the price-to-book ratio on the market today is about 4 to 1. If the market is willing to pay \$4.00 for the \$1.00 equity that is being reinvested, companies cannot be wasting the reinvested money.

SIEGEL: The confusing thing is that the price-to-book ratio for the S&P 500 or the DJIA is about 4 or 5 to 1 but the Tobin's *q*-ratio—which uses book value adjusted for inflation and replacement costs—is

nowhere near that amount. I think it could be very misleading to use historical market-to-book ratios.

LEIBOWITZ: Still, whether you use the market-to-book ratio or not, the idea of having high P/Es in an environment where monies are reinvested at less than the cost of capital produces the same inconsistency. Something doesn't compute.

IBBOTSON: The burden is on the people who are challenging the Miller-Modigliani theorem. M&M said that dividends and retention of earnings have the same effect so which number is used doesn't matter; you're saying it does matter.

ARNOTT: I believe the Miller-Modigliani theorem is an elegant formula that should work. But it doesn't match 130 years' worth of historical data.

IBBOTSON: We'll investigate that!

PHILIPS: In part, the difference may be something already mentioned: NIPA (which covers all businesses) versus the set of publicly traded securities (which is a subset of NIPA). Examining both groups separately might provide us some answers to the reinvestment question. Another angle on reinvestment is: Suppose we idealize the world so that businesses reinvest only what they need for their growth (so, it's a rational reinvestment, not empire building). What is our view now of how EPS should be growing? Is there a consensus? Rob Arnott has some very strong numbers showing that per share earnings grow more slowly than the economy. Will you be putting up that graph this afternoon, Rob?

ARNOTT: Yes, that's why I'm not saying anything.

SIEGEL: What's interesting is that growth has occurred over time in the marketable value of securities versus what would be implied by the NIPA profits. Many more companies are now public than used to be. A lot of partnerships have gone public in the

past 10–20 years. A lot of small companies, private companies, have gone public recently. Part of the reason could be the good stock market, and part could be a long-term trend. At any rate, in NIPA, a very big decline has occurred in "proprietors' income," which is derived from partnerships and individual owners, and an increase has occurred in corporate income as these private companies and partnerships went public. You have to be aware of this trend if you are using long-term data. It is one reason I think there is an upward trend in market value versus GDP. I'm not saying the ownership change alone explains the market value trend, or that it explains the whole amount, but changes between corporate income and noncorporate income are important.

IBBOTSON: So, as I've just said, either go to per share data to do this type of analysis or make sure you make all these adjustments to the aggregate data. See Diermeier, Ibbotson, and Siegel (1984) if you want to see how to make the adjustments.

TERHAAR: For the per share data, however, most people use the S&P 500, and the S&P 500 isn't really passive. It's a fairly actively managed index, particularly in recent years; the managers at Standard & Poor's have a habit of adding "hot" stocks, such as their July 2000 inclusion of JDS Uniphase. These substitutions have effects on the per share earnings and the growth rate that would not be present in a broader index or in the NIPA index.

SIEGEL: That's a very important point. Whenever the S&P 500 adds a company that has a higher P/E than the average company in the index, which has been very much the case in the past three years, the result is a dollar bias in the growth rate of earnings as the index is recomputed to make it continuous. My calculations show that the bias could be 1–2 percent a year in recent years as companies with extraordinarily high P/Es were added.

Current Estimates and Prospects for Change I

Robert J. Shiller

*Yale University
New Haven, Connecticut*

The equity premium puzzle and the foundations of behavioral finance are inseparable. The equity premium puzzle is a puzzle only if we assume that people's expectations are consistent with past historical averages, that expectations are rational. But behavioral finance has shown repeatedly the weakness of the assumption that rational expectations consistently drive financial markets. This presentation explores, in the context of recent stock market behavior, a number of reasons to doubt that rational expectations always find their way appropriately into stock prices. The reasons stressed have to do with psychological factors: (1) the difficulty that committees, groups, and bureaucracies have in changing direction, (2) the inordinate influence of the recent past on decisions, (3) the tendency (perhaps the need) to rely on "conventional wisdom," and (4) group pressure that keeps individuals from expressing dissent.

I will discuss here some issues in behavioral finance related to the so-called equity premium puzzle. The academic literature on the puzzle is based on the assumption that people are perfectly rational and consistent in their financial decision making and that their expectations for future returns are at all times in line with facts about past historical returns. The term "equity premium puzzle" refers to

the fact that the performance of the stock market in the United States has just been too strong relative to other assets to make sense from the standpoint of such rationality. But behavioral finance research has provided strong evidence against the very assumptions of rationality, at least against the idea that the rationality is consistent and responsive to relevant information and only relevant information. The equity premium puzzle and the foundations of behavioral finance are inseparable.

People's expectations cannot be equated with mathematical expectations, as the equity premium literature assumes. Expectations for future economic variables, to the extent that people even have expectations, are determined in a psychological nexus. I want to describe, in the context of recent experience in the stock market, some of the psychology that plays a role in forming these expectations. Considering recent experience will help provide concreteness to our treatment of expectations. The U.S. equity market became increasingly overpriced through the 1990s, reaching a phenomenal degree of overpricing by early 2000.¹ This event is a good case study for examining expectations in general.

I will be following here some arguments I presented in my 2000 book *Irrational Exuberance*, and I will also develop some themes that I covered in my 2002 paper, "Bubbles, Human Judgment, and Expert Opinion," which concentrated attention on the behavior of institutional investors—particularly, college endowment funds and nonprofit organizations (see Shiller 2002).

The theme of "Bubbles, Human Judgment, and Expert Opinion" is that even committees of experts can be grossly biased when it comes to actions like those that are taken in financial markets.

A lot of behavioral finance depicts rather stupid things going on in the market, but (presumably) trustees and endowment managers are pretty intelligent people. Yet, they, as a group, have not been

¹ See the testimony by John Y. Campbell and Robert J. Shiller before the Federal Reserve Board on December 3, 1996. Summarized in Campbell and Shiller (1998).