**Valuation Ratios and the Long-Run Stock Market Outlook:
An Update**

By

John Y. Campbell and Robert J. Shiller

March 2001

COWLES FOUNDATION DISCUSSION PAPER NO. 1295

**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

Valuation Ratios and the Long-Run Stock Market Outlook: An Update¹

John Y. Campbell and Robert J. Shiller²

Abstract. The use of price–earnings ratios and dividend–price ratios as forecasting variables for the stock market is examined using aggregate annual US data 1871 to 2000 and aggregate quarterly data for twelve countries since 1970. Various simple efficient-markets models of financial markets imply that these ratios should be useful in forecasting future dividend growth, future earnings growth, or future productivity growth. We conclude that, overall, the ratios do poorly in forecasting any of these. Rather, the ratios appear to be useful primarily in forecasting future stock price changes, contrary to the simple efficient-markets models. This paper is an update of our earlier paper (1998), to take account of the remarkable behavior of the stock market in the closing years of the twentieth century.

Keywords: Stock market, price–earnings ratio, forecasts, expectations, dividend–price ratio, efficient markets, Standard & Poor's 500, present value, productivity.

JEL Classification: G12

When stock market valuation ratios are at extreme levels by historical standards, as dividend–price and price–earnings ratios have been for some years in the US, one naturally wonders what this means for the stock market outlook. It seems reasonable to suspect that prices are not likely ever to drift too far from their normal levels relative to indicators of fundamental value, such as dividends or earnings. Thus it seems natural to give at least some weight to the simple mean-reversion theory that when stock prices are very high relative to these indicators, as they have been recently, then prices will eventually fall in the future to bring the ratios back to more normal historical levels. The

¹ This article is based on our joint testimony before the Board of Governors of the Federal Reserve System, December 3, 1996, on material circulated in Shiller (1996) and on Campbell and Shiller (1998).

² Department of Economics, Harvard University, Littauer Center 213, Cambridge, MA 02138, 617-496-6448, john_campbell@harvard.edu; and Cowles Foundation for Research in Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06520, 203-432-3708, shiller@econ.yale.edu. We acknowledge the able research assistance of Elena Rangelova and Daniel Waldman, help with data from Robert J. Gordon, and the helpful comments of Paul Samuelson.

idea that they should do so seems intuitive and basic. Metaphorically, when one is mountaineering, one can enjoy the exhilarating view from high up on a mountain, and may look forward to the possibility of discovering a way up to a much higher level. But one will reflect that, realistically, at a random date years from now, one will probably be back down at ground level.

On December 3, 1996, we testified before the Federal Reserve Board that, despite all the evidence that stock returns are hard to forecast in the short run, this simple theory of mean reversion is basically right and does indeed imply a poor long-run stock market outlook. We amplified our testimony and published it in 1998, continuing to assert our pessimistic long-run scenario.³

The stock market did not immediately move to encourage faith in our theory. Since our testimony, the stock market, as measured by the real (inflation-corrected) Standard & Poor Composite index, has increased by 80% above its value when we testified, and 30% above its value when we published.

Despite these developments, we believe that our original testimony and article are even more relevant today. Valuation ratios moved up in the year 2000 to levels that were absolutely unprecedented, and are still nearly as high as of this writing at the beginning of 2001. Even allowing for the possibility that the economy and financial markets have undergone some structural changes, these ratios imply a stronger case for a poor stock market outlook than has ever been seen before. To underscore this conviction, we present here an extended version of our 1998 paper, with data updated to 2000.

I. Historical Behavior of Valuation Ratios

We should first understand what the stability of a valuation ratio itself implies about mean reversion. If we accept the premise for the moment that valuation ratios will continue to fluctuate within their historical ranges in the future, and neither move permanently outside nor get stuck at one extreme of their historical ranges, then when a valuation ratio is at an extreme level either the numerator or the denominator of the ratio must move in a direction that restores the ratio to a more normal level. *Something* must be forecastable based on the ratio, either the numerator or the denominator. For example, high prices relative to dividends — a low dividend–price ratio — must forecast some combination of unusual increases in dividends and declines (or at least unusually slow growth) in prices.

The conventional random-walk theory of the stock market is that stock price changes are not predictable, so that neither the dividend–price ratio nor any other valuation ratio has any ability to forecast movements in stock prices. But then, if the random-walk theory is not to imply that the dividend–price ratio will move beyond its historical

³ Over this interval we have also published related papers, Campbell (1999) and Shiller (1999), and have written two books, Campbell and Viceira (2001) and Shiller (2000), that expand on our views.

range or get stuck forever at the current extreme, it requires that the dividend-price ratio predicts future growth in dividends.⁴

Does the dividend-price ratio forecast future dividend movements as required by the random-walk theory, or does it instead forecast future movements in stock prices? We answer this question using a long-run annual US data set that extends today's S&P 500 Index back in time to 1872.⁵ The answer is given by the pair of scatterplots shown in Figure 1. Each scatterplot has the dividend-price ratio, measured as the previous year's dividend divided by the January stock price, on the horizontal axis. (The horizontal axis scale is logarithmic but the axis is labeled in levels for ease of reference.) Over this period the historical mean value for the dividend-price ratio was 4.65%.

In the top part of the figure the vertical axis is the growth rate of real dividends (measured logarithmically as the change in the natural log of real dividends) over a time interval sufficient to bring the dividend-price ratio back to its historical mean of 4.65%. More precisely, we measure the dividend growth rate from the year preceding the year shown until the year before the dividend-price ratio again crossed 4.65%. Because dividends enter the dividend-price ratio with a one-year lag, this is the appropriate way to measure growth in dividends from the base level embodied in a given year's dividend-price ratio to the level that prevailed when the dividend-price ratio next crossed its historical mean.⁶

⁴ The random-walk theory is a special case of the efficient-markets theory of stock prices. In general, the efficient-markets theory allows the equilibrium rate of return required by investors to vary over time. (See for example Campbell and Cochrane, 1999.) The random-walk theory assumes that this required rate of return is constant. We are in fact oversimplifying the random-walk theory in this paper, because the theory actually says that stock *returns*, not prices, should be unforecastable. Since the dividend-price ratio is itself a component of the stock return, the random-walk theory says that a lower dividend-price ratio should be associated with slightly more rapid price growth to offset the lower dividend component of return. In other words, the theory says that prices should move in a direction that drives the dividend-price ratio away from its historical average; dividends must do more than all the adjustment necessary to bring the ratio back to its historical average. However, the difference between return and price change is small and in practice forecasts of returns and forecasts of price changes are very similar. See our (1988a) paper for a careful analysis of dividend forecasts within the context of a log-linearized mathematical representation of the efficient-markets theory, or Campbell, Lo, and MacKinlay (1997), Chapter 7, for a recent textbook exposition.

⁵ The data in this paper use the January Standard and Poor Composite stock price for each year since 1872, while earnings and dividends are for the entire previous year. Data before 1926 are based on Cowles (1939). The price index used to deflate nominal values to real values is the producer price index. See Shiller (1989) for a description of these data.

⁶ The time intervals required to bring the dividend-price ratio back to its mean typically exceed one year, so the dividend growth rate for any particular year can affect several successive observations. This overlapping of successive time intervals implies that the different points in the scatterplot are not statistically independent. There are, however, 29 nonoverlapping time intervals in our sample, so the data are not insubstantial. Statistical tests of the significance of analogous relations with fixed

Since 1872, the dividend–price ratio has crossed its mean value 29 times, with intervals between crossings ranging from one year to twenty years (the twenty-year interval being between 1955 and 1975). The different years are indicated on the scatter diagram by two-digit numbers; a * after a number denotes a 19th Century date. The last year shown is 1983, since this is the last year that was followed by the dividend–price ratio crossing its mean. (The ratio has been below its mean ever since.) A regression line is fit through these data points, and a vertical line is drawn to indicate the dividend–price ratio at the start of the year 2000. The implied forecast for dividend growth is the horizontal dashed line marked where the vertical line intersects the regression line.

It is obvious from the top part of Figure 1 that the dividend–price ratio has done a poor job as a forecaster of future dividend growth to the date when the ratio is again borne back to its mean value. The regression line is nearly horizontal, implying that the forecast for future dividend growth is almost the same regardless of the dividend–price ratio. The R^2 statistic for the regression is 0.25%, indicating that only one-quarter of one percent of the variation of dividend growth is explained by the initial dividend–price ratio.

It must follow, therefore, that the dividend–price ratio forecasts movements in its denominator, the stock price, and that it is the stock price that has moved to restore the ratio to its mean value. In the lower part of Figure 1 the vertical axis shows the growth rate of real stock prices (measured logarithmically as the change in log real stock prices) between the year shown and the next year when the dividend–price ratio crossed its mean value. The scatterplot shows a strong tendency for the dividend–price ratio to predict future price changes. The regression line has a strongly positive slope, and the R^2 statistic for the regression is 63%. We have answered our question: It is the denominator of the dividend–price ratio that brings the ratio back to its mean, not the numerator.

At the start of 2000, the dividend–price ratio was only 1.2%, well to the left of any points shown in the figure. The lower part of Figure 1 shows that on previous occasions when the dividend–price ratio has been below 3.4%, the stock market has *always* declined in real terms over the interval to the next crossing of the mean dividend–price ratio; real declines in stock prices have *always* played a role in restoring such extreme low dividend–price ratios to the mean. The fitted value of the regression line for 2000 indicates that the next time that the dividend–price ratio is back to its mean, the log real value of the stock market will be more than 1.6 lower than it is today. Translating into percentage terms, this says that the stock market will lose more than three-quarters of its real value! Can we take such a forecast seriously? What modifications should we make to such a forecast?

horizons, taking account of the overlapping intervals, are reported in Campbell and Shiller (1988b, 1989).

Fixed-horizon forecasts from the dividend–price ratio

Figure 1 shows the powerful ability of the dividend–price ratio to predict price movements to the date at which the dividend–price ratio next crosses its mean. We looked at the figure to see what it is that restores the ratio to its mean: the numerator or the denominator. But, the problem with these forecasts is that we do not know when the dividend–price ratio will next cross its mean; historically this has ranged from one to twenty years. We now show scatterplots like Figure 1, but where the vertical axis is changed to show growth rates of dividends and prices over a fixed horizon. The horizon is one year in Figure 2, and ten years in Figure 3. We should expect to see a worse fit than in Figure 1, of course, since with these figures we do not measure dividend and price growth rates over intervals when the ratio returned to its mean value.

The upper part of Figure 2 shows that over one year, the dividend–price ratio does forecast dividend growth with the negative sign predicted by the efficient-markets theory. Years in which January stock prices are high, relative to last year’s dividends, tend to be years in which this year’s dividends are high relative to last year’s dividends. The dividend–price ratio is able to explain 13% of the annual variation in dividend growth. Such short-horizon forecasting power should not be surprising; dividends are fairly predictable over a few quarters, and the January stock price is measured well after most of last year’s dividends have been paid, at a time when it may be relatively easy for market participants to anticipate the level of dividends during the coming year.

The lower part of Figure 2 shows that the dividend–price ratio has little forecasting power for stock price changes over the next year. Prices do have a very slight tendency to fall in years when they are initially high relative to dividends, but this relationship explains less than 1% of the annual variance of stock prices. The short-run noise in stock prices swamps the predictable variation that was visible in Figure 1.⁷

In Figure 3, however, where the horizon is ten years rather than one year, many of the patterns of Figure 1 become apparent again. Just as in Figure 1, there is only a very weak relation between the dividend–price ratio and subsequent ten-year dividend growth. In fact the relation in Figure 3 is even less consistent with the efficient-markets theory than the relation in Figure 1, because the Figure 3 relation is positive, implying that dividends tend to move in the wrong direction to restore the dividend–price ratio to its historical average level. Just as in Figure 1, there is a substantial positive relation between the dividend–price ratio and subsequent ten-year price growth. The R^2 statistics are a trivial 1% for dividend growth but 9% for price growth.

The unusual recent behavior of the stock market is visible in the bottom panels of both Figures 2 and 3. In Figure 2, the low dividend–price ratios and large price increases of the years 1995 through 1999 are visible as five points at the top left of the figure. In Figure 3, price increases during the 1990s have a somewhat smaller effect but are visible in three points for the years 1988, 1989, and 1990 at the top left of the figure.

⁷ Campbell, Lo, and MacKinlay (1997), Chapter 7, explains in more formal terms how R^2 statistics can rise with the length of the horizon over which returns are measured.

Given the low value for the dividend–price ratio at the start of 2000, the regression in the bottom panel of Figure 3 implies a decline of 0.6 in the log real stock price over the next ten years. This corresponds to a 55% loss of real value.⁸

Alternative valuation ratios

The dividend–price ratio is a widely used valuation ratio, but it has the disadvantage that its behavior can be affected by shifts in corporate financial policy, a point we discuss later in the paper. Accordingly it is worthwhile to explore alternative measures of the level of stock prices.

Figure 4 illustrates some key valuation ratios in our long-run annual US data set. The top left panel of the figure shows the price–earnings ratio, calculated using the January stock price in each year divided by the level of earnings from the previous year. The bottom left panel shows the dividend–price ratio, calculated using the dividends from the previous year divided by the January stock price. These ratios are not adjusted to express them in real terms, because it is assumed that the same general price index applies to the earnings or dividend series and the stock price series.

Figure 4 illustrates the fact that price–earnings ratios have normally moved in a range from 8 to about 20, with a mean of 14.5 and occasional spikes down as far as 6 or up as high as 26. At the beginning of 2000, the price–earnings ratio was high, at 29.6, but not at a record level. Dividend–price ratios have normally moved in a range from 3% to about 7%, with a mean of 4.65% and occasional movements up to almost 10%. Very recently the dividend–price ratio has fallen to a record low of 1.2%, well below the historical range.

Since stock price increases drive up price–earnings ratios and drive down dividend–price ratios, it is not surprising that the two series in Figure 4 generally move opposite to one another. There are, however, various spikes in the price–earnings ratio that do not show up in the dividend–price ratio. These spikes occur when recessions temporarily depress corporate earnings. Since we use previous-year earnings to calculate price–earnings ratios, depressed earnings in 1921, 1933, and 1991, for example, show up in our price–earnings series in 1922, 1934, and 1992.

A clearer picture of stock market variation emerges if one averages earnings over several years. Benjamin Graham and David Dodd, in their now famous 1934 textbook *Security Analysis*, said that for purposes of examining valuation ratios, one should use an average of earnings of “not less than five years, preferably seven or ten years” (p. 452).

⁸ As we mentioned in footnote 3, stock returns differ from stock price changes because they include the direct contribution of dividends. Figure 3 implies an unusually poor year-2000 outlook for stock returns, for three reasons. First, dividends are initially low relative to prices. Second, the top part of Figure 3 shows that dividends are predicted to grow slowly over the next ten years. Third, the bottom part of Figure 3 shows that real prices are predicted to fall over the next ten years. A scatterplot with ten-year real stock returns on the vertical axis looks much like the bottom part of Figure 3, but with a better fit (an R^2 statistic of 16% rather than 9%). The cumulative continuously compounded ten-year return forecast implied by the January 2000 dividend-price ratio is –44%.

Following their advice we smooth earnings by taking an average of real earnings over the past ten years.⁹ The top right panel of Figure 4 shows the ratio of the January real stock price to smoothed real earnings from the previous year. This price-smoothed-earnings ratio responds to long-run variations in the level of stock prices. It has roughly the same range of variation as the conventional price-earnings ratio, with a slightly higher mean of 16.0, but the record high of 44.9 now appears at the start of 2000. This record ratio dwarfs the previous record of 28.0, set in 1929.

The bottom right panel of Figure 4 shows the ratio of current real earnings to smoothed real earnings. This figure shows that in 2000 real earnings have indeed grown quite well when compared to their ten-year past average, but this earnings growth is not record breaking and there are a number of comparable experiences in history. It is price growth, not earnings growth, that has set all-time records lately.

Forecasts from the price-smoothed-earnings ratio

Figures 5 and 6 have the same format as Figures 2 and 3, except that the ratio of price to a ten-year moving average of real earnings appears on the horizontal axis of each scatterplot, and we look at the growth rate of the ten-year moving average of earnings rather than the growth rate of dividends. The price-smoothed-earnings ratio has little ability to predict future growth in smoothed earnings; the R^2 statistics are 1% over one year and 5% over ten years. However, the ratio is a good forecaster of ten-year growth in stock prices, with an R^2 statistic of 30%. The fit of this relation is substantially better than we found for the dividend-price ratio in Figure 3.¹⁰

Noting that the price-smoothed-earnings ratio for January 2000 is a record 44.9, the regression illustrated in Figure 6 is predicting a catastrophic ten-year decline in the log real stock price. We do not find this extreme forecast credible; when the independent variable has moved so far from the historically observed range, we cannot trust a linear regression line. However, this extreme forecast does, we think, suggest some real concerns that future price growth will be small or negative.

⁹ We first looked at smoothed real earnings in our (1988b) paper. There we averaged log real earnings rather than levels of real earnings (that is, we used a geometric rather than an arithmetic average), but this makes little difference to the results. We also compared ten-year and thirty-year moving averages of earnings, and found that they have similar properties.

¹⁰ The price-smoothed-earnings ratio is also a much better predictor than the conventional price-earnings ratio. The noise in annual earnings distorts the fundamental relation illustrated in Figure 6. The superior forecasting power of the price-smoothed-earnings ratio carries over to ten-year real returns; a regression of ten-year returns on the price-smoothed-earnings ratio has an R^2 statistic of 40%, whereas a regression of ten-year returns on the dividend-price ratio has an R^2 statistic of only 16%.

Ratios' forecasts of productivity

Popular commentators on the stock market often justify high valuation ratios by reference to expectations of future productivity growth, that is, future growth in output per manhour, as if productivity were another indicator of the value of firms. They point to rapid productivity growth in the second half of the 1990s and argue that the stock market rationally anticipates a continuation, or even an acceleration, of this trend.¹¹ A difficulty with this line of argument is that higher output per manhour in the future may well accrue to workers, or to the entrepreneurs who create new firms, rather than to the owners of existing firms. Nonetheless it is interesting to ask whether the stock market has historically predicted variations in productivity growth. We can extend our previous analysis by substituting productivity growth, in place of earnings growth, as the variable to be forecasted.

The top panel of Figure 7 shows the log of real output per hour for the nonfarm housing private economy, along with the same log real earnings series we plotted in Figure 4.¹² Note that productivity has virtually the same growth rate as real S&P earnings, but productivity has much less volatility, it more nearly hugs a trend line. The bottom panel of Figure 7 shows ten-year growth rates of output per hour and real earnings. There are some short-run comovements in the two series, probably reflecting the short-run effects of recessions on both profits and productivity.

Figure 8 is a scatter diagram with the ratio of price to smoothed ten-year earnings on the horizontal axis, and the subsequent ten-year growth in productivity on the vertical axis. We see that the price-smoothed-earnings ratio has virtually no ability to predict future productivity growth. A scatter with the conventional price-earnings ratio on the horizontal axis, not shown here, is even less successful. These results do not support the view that movements in stock prices reflect rational forecasts of future productivity growth.

II. Is The Twenty-First Century A New Era?

Over the past century the American economy has been transformed in many fundamental ways. Agriculture gave way to industry, and industry has given way to services as the economy's leading sector. Automobiles and airplanes have revolutionized transport, while radio, television, and now the Internet have transformed communication. Massive corporations emerged to exploit the economies of mass production, but these are now being replaced by smaller, more flexible organizations that can exploit information technology more effectively.

¹¹ While popular discussion normally emphasizes productivity growth in a few new economy sectors, Nordhaus (2000) shows that productivity growth rate increases in the 1990s were widespread and not narrowly focussed on these sectors.

¹² Robert J. Gordon (2000, Figure 3, p. 66) supplied the productivity series. His multifactor productivity is fairly similar in appearance to the output per hour series used here.

These changes have affected the financial sector just as deeply as any other part of the economy. Yet certain aspects of financial market behavior have remained remarkably stable throughout the tumult of the 20th Century. We have seen that stock market valuation ratios have moved up and down within a fairly well-defined range, without strong trends or sudden breaks.

Despite the historical stability of valuation ratios, some market observers question whether historical patterns offer a reliable guide to the future. Various arguments are put forward to justify the notion that financial markets are entering a “new era.” Some of these arguments have to do with corporate financial policy, while others concern investor behavior or the structure of the US economy. We now briefly review some of these arguments.

Repurchases and the dividend-price ratio

Dividends represent cash paid to ongoing shareholders, and this makes dividends an appealing indicator of fundamental value. In fact, over very long holding periods the return to shareholders is dominated by dividends, because the end-of-holding-period stock price becomes trivially small when it is discounted from the end to the beginning of a long holding period.

Nonetheless, an important criticism of the dividend-price ratio is that it can be affected by corporate financial policy. As a tax-favored alternative to paying dividends, companies can repurchase their stock. Repurchases transfer cash to those shareholders who sell their stock, and benefit ongoing shareholders because future dividend payments will be divided among fewer shares. If a corporation permanently diverts funds from dividends to a repurchase program, it reduces current dividends but begins an ongoing reduction in the number of shares and thus increases the long-run growth rate of dividends per share. This in turn can permanently lower the dividend-price ratio, driving it outside its normal historical range. Many commentators have argued that repurchases, not excessive stock prices, are responsible for record low dividend-price ratios in the late 1990s.

One way to adjust the dividend-price ratio for shifts in corporate financial policy is to add net repurchases (dollars spent on repurchases less dollars received from new issues) to dividends. Cole, Helwege, and Laster (1996) did this for S&P 500 firms over the period 1975–1996 and found that dividend-price ratios should be adjusted upwards significantly during the mid-1980s and the mid-1990s, for example by 0.8% in 1996. This approach assumes that both repurchases and issues of shares take place at market value, so that dollars spent and received correspond directly to shares repurchased and issued. In practice, however, many companies issue shares below market value as part of their employee stock option incentive plans. Liang and Sharpe (1999) correct for this in a study of the largest 144 firms in the S&P 500; they find that the dividend-price ratio for those firms should be adjusted upwards by 1.39% in 1997 (a number that they argue is not sustainable in the long run) and 0.75% in 1998.

A glance at Figure 4 shows that an adjustment of this magnitude brings the dividend-price ratio back closer to the bottom of its normal historical range, but does not

bring it anywhere close to the middle of the normal range. For this reason, and because repurchase programs do not affect price–earnings ratios, corporate financial policy cannot be the only explanation of the abnormal valuation ratios observed in recent years.

Intangible investment and the price–earnings ratio

A criticism that is commonly directed against use of the conventional price–earnings ratio as an indicator of stock market valuation is that the denominator of the ratio, earnings, has become biased downward because the new economy involves substantial investments in intangibles, which are, following conventional accounting procedures, deducted from earnings as current expenses. For example, it is a hallmark of many companies in the new economy that they plan to attract a large volume of customers but to lose money for years, hoping that the high level of activity will enable them to build an effective high-tech business organization as well as to solidify public acceptance of their product. The cost of activities that promote such intangible capital, these critics argue, should not really be deducted directly from earnings, since they are effectively long-term investments.

Hall (2000) has called such intangible capital “e-capital,” and argues that there has been a great deal of investment in e-capital in the 1990s “resulting at least in part from technological progress in forming e-capital.”¹³ He tells a story in which the accumulation of e-capital in the 1990s explains the much higher measured price–earnings ratios as well as the higher wages paid to college graduates (who create e-capital) relative to the wages paid to those who did not graduate from college. His story is also consistent with the fact that gains in measured productivity in the late 1990s have been modest, well below the gains of the 1950s and 1960s.

McGrattan and Prescott (2001) present corrections that should be made to earnings due to investment in intangible capital, and also attempt to gauge the effects through time of changes in tax rates on the price/earnings ratio. They conclude, in the context of a complete calibrated general equilibrium model, that the price earnings ratio 1987-99 should be roughly twice what it was in 1955-62. The main reason that the price–earnings ratio should have doubled over that time interval in their model is that they estimate that the personal income tax rate on dividends has fallen sharply over that interval. Their calibrated model represents that that tax rate has fallen from 50% to 9% over that time interval, reflecting the advent since 1962 of 401(K) and other tax-sheltered investment vehicles.¹⁴

These new economy stories are interesting possibilities to explain the stock market, but they are just stories: no convincing justification has been given for assuming that investment in intangibles is really dramatically more important in recent years than it was in earlier years. Neither Hall nor McGrattan and Prescott show that their models fit long

¹³ Hall “E-Capital...” (2000, p. 76). See also Hall (2001).

¹⁴ See McGrattan and Prescott (2001).

historical time series data. Their calibrated models explain only recent observations, and hence their fit is of little persuasiveness in judging the current valuation ratios. Hall's model would imply that most US firms had negative e-capital in the period 1980–87.¹⁵

Bond and Cummins (2000), using data on 459 individual firms over the period 1982–1998, partially measure intangible capital investment by expenditures the firms make on research and development and on marketing. They consider the effect of intangible capital on investment equations and conclude that intangible investments do not appear to justify the current high valuation in the market.¹⁶

The baby boom, market participation, and the demand for stock

Many observers suggest that there has been a secular shift in the attitudes of the investing public towards the stock market. As the baby-boom generation comes to dominate the economically and financially active population, its attitudes become more important while those of earlier generations have less and less weight. It is argued that baby-boomers are more risk-tolerant (perhaps because they do not remember the extreme economic conditions of the 1930s), and that they tend to favor stocks over bonds (perhaps because they are influenced by the extremely poor performance of bonds during the inflationary 1970s). Thus valuation ratios may be extreme today because baby-boomers are willing to pay high prices for stocks; the ratios may remain extreme for as long as this demographic effect persists — that is, well into the 21st Century — and may even move further outside their historical ranges if the demographic effect strengthens.

A variant of this argument emphasizes that economists have had great difficulty in reconciling historical stock price levels with standard equilibrium asset pricing models. Mehra and Prescott (1985) pointed out that stock prices have been much lower than standard models would predict; they initiated an enormous literature on the “equity premium puzzle,” but no entirely convincing explanation has been found. Perhaps the baby-boom generation is the first to realize that historical valuation ratios were a mistake, and recent stock price movements represent a correction of the mistake.¹⁷

¹⁵ Jason Cummins (2000), p. 109).

¹⁶ Bond and Cummins (2000) find that the coefficient of Tobin's Q in the investment equation is usually not significant after including their additional intangible capital measures unless one replaces the usual Tobin's Q ratio with a ratio that has as its numerator not the actual price but the forecasted present value of future earnings. Thus, firms' behavior suggests that managers themselves do not believe the valuations for the market even after partial correction for intangible investment.

¹⁷ Siegel (1994) presents a moderate version of this argument, while Glassman and Hassett (1999) present an extreme version, arguing that stocks should not carry any risk premium at all, and that stock prices will rise dramatically further once investors come to realize this fact. Both Siegel and Glassman–Hassett emphasize that stock returns have historically had lower risk at long horizons than at short horizons. This is a manifestation of the same mean-reversion

Alternatively, baby-boomers may benefit from institutional innovations that make it easier for less well-off people to participate in the stock market, and to hold diversified portfolios. Heaton and Lucas (1999) and Vissing-Jørgenson (1998) show that broader participation and cheaper diversification can drive up the demand for stock and increase stock prices. However such effects are unlikely to explain large movements in the stock market because most wealth is now, and always has been, controlled by wealthy people who face few barriers to stock market participation and diversification.

In support of this line of thought, it has been pointed out that the dividend-price ratio shows some evidence of trend decline during the whole of the period since World War II. The appearance of long-run stability in this ratio in Figure 4 would be much weaker if the figure began in the mid-20th Century rather than in 1872.¹⁸ On the other hand long-run trends in stock-market participation are not plausible candidates to explain the sharp run-up in stock prices during the late 1990s.

While it may be true that the demand for stock has increased, this does not necessarily contradict the pessimistic stock market outlook presented earlier in this paper. The argument is that demand has driven stock prices up relative to dividends and earnings. But since the demand for stock does not change the expected paths of future dividends and earnings, higher stock prices today must depress subsequent stock returns unless demand is even stronger at the end of the holding period. Over the ten-year holding period emphasized in this paper, there does not seem to be any good reason to expect stock demand to strengthen further from today's high levels.

Also, it may not be correct to think of investors' attitudes as shifting only slowly, in reaction to long-run demographic changes. Economic conditions may also be important. It is noticeable that stock prices tend to be high relative to indicators of fundamental value at times when the economy has been growing strongly. This tendency is visible in Figure 4; high price-earnings and price-smoothed-earnings ratios and low dividend-price ratios are characteristic of periods, such as the 1920s, 1960s, and mid-1990s, when real earnings have been growing rapidly so that current earnings are well above smoothed earnings. If economic growth in general, or earnings growth in particular, influence investors' attitudes then weaker economic conditions could rapidly bring prices back down to more normal levels.¹⁹

documented in this paper, but Siegel and Glassman-Hassett do not stress the low return forecasts that are implied by mean-reversion.

¹⁸ Blanchard (1993) emphasizes the trend postwar decline in the dividend-price ratio and in various other measures of the risk premium investors demand for holding stocks. Bakshi and Chen (1994) argue that demographic effects can explain the high stock market of the 1960s and 1980s and low stock market of the 1970s, but they do not ask whether their demographic measures have explanatory power for other countries or time periods.

¹⁹ This leaves open the question of why investors' attitudes might be affected by economic conditions. Barsky and De Long (1993) and Barberis, Shleifer, and Vishny (1998) argue that investors irrationally extrapolate recent earnings growth into the future, so that the stock market becomes overvalued when earnings growth has been strong. Campbell and Cochrane (1999) argue

Inflation

Other observers have argued that today's high stock prices can be justified by the steady decline in inflation that has taken place since the early 1980s. These observers point out that since 1960, the dividend–price ratio has moved closely with the inflation rate and with the yield on long-term government bonds, which is closely associated with expectations of future inflation. Thus it should not be surprising to see high stock prices given low recent inflation.

There are two weaknesses in this argument. First, the correlation between stock prices and inflation is much stronger before the mid-1990s than during the last five years. It is hard to explain the recent rise in the stock market by any large change in the inflation outlook.

Second, it is not clear that the association between stock prices and inflation is consistent with the efficient-markets theory that stock prices reflect future real dividends, discounted at a constant real interest rate. That is, low inflation may help to *explain* high stock prices but may not *justify* these prices as rational. Modigliani and Cohn (1979) argued twenty years ago that the stock market irrationally discounts real dividends at nominal interest rates, undervaluing stocks when inflation is high and overvaluing them when inflation is low.²⁰ At that time their argument implied stock market undervaluation; today the same argument would imply overvaluation. Whether or not one accepts Modigliani and Cohn's behavioral hypothesis, it should be clear that the relation between inflation and stock prices does not necessarily contradict our pessimistic long-run forecast for stock returns.

III. International Evidence

We have emphasized that in US data prices, rather than dividends or earnings, appear to adjust to bring abnormal valuation ratios back to historical average levels. Do other countries' stock markets behave in the same way, or is the US experience anomalous?

Unfortunately very little long-term data are available for most stock markets. One standard data source is Morgan Stanley Capital International, but these data go back only to 1970 or so. To appreciate how short this sample is, note from Figure 4 that since the early 1970s the time-series plot of the US dividend–price ratio has been dominated by a single hump-shaped pattern. With under thirty years of data, it is not sensible to use a ten-year horizon, so we reduce the horizon to four years.

Figure 9 presents scatterplots like Figures 2 and 3, but with quarterly data and a four-year horizon. The dividend–price ratio appears on the horizontal axis of each

that investors become more risk-tolerant when the economy is strong, because their well being is determined by their consumption relative to past standards, rather than by the absolute level of consumption.

²⁰ Ritter and Warr (1999) have recently revisited this hypothesis.

scatterplot, and four-year dividend or price growth appears on the vertical axis. The first quarter of each year is indicated with a year number; the other quarters are marked with a cross. Results are shown for twelve countries: Australia, Canada, France, Germany, Italy, Japan, the Netherlands, Spain, Sweden, Switzerland, the United Kingdom, and, for comparison, the United States.²¹

The countries in Figure 8 fall into three main groups. The English-speaking countries, Australia, Canada, and the UK, behaved over this short sample period very much like the US. The dividend–price ratio was positively associated with subsequent price growth, and showed little relation to subsequent dividend growth. Several Continental European countries, France, Germany, Italy, Sweden, and Switzerland, showed a very different pattern over this sample period. In these countries a high dividend–price ratio was associated with weak subsequent dividend growth, just as the efficient-markets theory would imply. There was little relation between the dividend–price ratio and subsequent price growth. Japan and Spain represent an intermediate case in which the dividend–price ratio appears to have been associated with both subsequent dividend growth and subsequent price growth. Finally the Netherlands show no clear relation between the dividend–price ratio and subsequent growth rates of either dividends or prices.

These recent international data provide mixed evidence. Recent price movements, often the very price movements that have made the valuation ratios so anomalous today, have a large effect on the scatterplots in Figure 9, and this makes them somewhat hard to interpret.

IV. Some Statistical Pitfalls

Some subtle statistical issues arise when one tries to draw conclusions from scatter diagrams such as those presented here. Since the observations are overlapping whenever the horizon is greater than one year (or one quarter in Figure 9), the different points are not statistically independent of one another. We must correct for this problem in judging the statistical significance of our results. Also, valuation ratios are random rather than deterministic, and it is well known that regressions with random regressors can have biased coefficients in small samples.

Let us consider the conclusions that we drew from looking at Figure 1. We noted that the slope of the regression line in the top part of the figure, predicting log real dividend growth over the time interval to the next crossing of the mean of the dividend–price ratio, was not substantially negative as the efficient-markets theory would predict. Were we right to conclude that real dividends do not behave in accordance with the efficient-markets theory? Or are our regression results possibly spurious?

In the 1998 version of this paper we did a simple Monte Carlo experiment to study this issue. We constructed artificial data in which the dividend–price ratio does not forecast future price changes over any fixed horizon. In other words, we generated data

²¹ See Campbell (1998) for a more detailed analysis of these international data.

that satisfy the efficient-markets prediction that the real stock price is a random walk.²² Also, we generated the data to match several important characteristics of the actual annual US data.

We began by estimating a first-order autoregressive (AR(1)) model for the log dividend–price ratio using our 125 observations for the period 1872 to 1997. We corrected the regression coefficient for small-sample bias using the Kendall correction, obtaining a coefficient of 0.81. Using a random normal number generator with the estimated standard error of the error term in the bias-corrected regression, and using a random normal starting value whose variance equals the unconditional variance for this AR(1) model, we generated 125 observations of a simulated AR(1) log dividend–price ratio. Next, we generated 125 observations of a simulated random walk for the log real stock price, using a random normal number generator with the estimated standard deviation of the actual change in the log real price. In the actual data, changes in the stock price and in the dividend–price ratio have a negative covariance; we also matched this covariance in our artificial data. Finally, we generated a log real previous-year dividend by adding the log dividend–price ratio and the log stock price.

We repeated this exercise 100,000 times. In each iteration, we used the artificial data to produce scatters and regression lines based on 125 observations like those shown in the top part of Figure 1. We found that the average number of crossings of the mean of the dividend–price ratio was 26.5, not far from the number of 29 observed with our actual data. But in 100,000 iterations we found that the slope of the regression line shown in the top part of Figure 1 was almost always much more negative than the estimated slope with the actual data. The estimated slope in the artificial data was greater than the estimated slope with actual data (–0.04) only 0.02%, two hundredths of one percent, of the time. The estimated regression coefficient in these Monte Carlo iterations tended to be close to minus one, very far from the almost-zero slope coefficient represented by the line in the figure. In this respect, our Monte Carlo results are extremely different from the results with the actual data. We conclude that our result in the top part of Figure 1 is indeed anomalous from the standpoint of the efficient-markets theory.

Next we used the change in the log real stock price as the dependent variable in the Monte Carlo experiment, so that in each iteration we estimated the regression line shown in the bottom part of Figure 1. In 100,000 iterations we never once obtained a regression coefficient as great as the slope coefficient of 1.25 shown in the bottom part of Figure 1. While the average estimated slope coefficient in the Monte Carlo experiments is

²² As before, we are oversimplifying the efficient-markets theory by ignoring the distinction between price changes and returns. In Campbell and Shiller (1989) we generated artificial data for a Monte Carlo study in which returns, rather than stock price changes, are unpredictable. This procedure is considerably more complicated, however, and it only makes the patterns seen in the actual data more anomalous.

positive, the average value is only 0.18, far below the estimated coefficient with actual data.²³

Other Monte Carlo experiments relevant to judging the results in this paper are reported in Campbell and Shiller (1989), Goetzmann and Jorion (1993), Nelson and Kim (1993), and Kirby (1997). Nelson and Kim (1993) generate artificial data from vector autoregressions (VARs) of stock returns and dividend yields on lagged returns and yields. The artificial stock-return series are constructed to be unforecastable but correlated with innovations in dividend yields. Campbell and Shiller (1989) follow a similar approach.²⁴ Nelson and Kim find that 10-year regression coefficients and R^2 statistics are highly unlikely to be as large as those found in the actual data if expected stock returns are truly constant. Campbell and Shiller's results are consistent with this finding.

Goetzmann and Jorion (1993) use a different approach. They construct artificial data using randomly generated returns and historical dividends, which of course are fixed across different Monte Carlo runs. They combine these two series to get random paths for dividend yields. The problem with this methodology is that it produces nonstationary dividend yields that have no tendency to return to historical average levels. Thus Goetzmann and Jorion avoid the need for dividend yields to forecast either dividend growth or price growth; in their simulations stock prices are equally uninformative about fundamental value and about future returns. Goetzmann and Jorion also confine their attention to horizons of four years or less. Large long-horizon regression coefficients and R^2 statistics occur somewhat more often in the Goetzmann–Jorion Monte Carlo study than in the Nelson–Kim study, but the four-year results in the actual data remain quite anomalous.

Kirby (1997) uses Monte Carlo methods to further illustrate biases that can arise in conventional statistical tests of market efficiency. Kirby's results are not very relevant to our regressions, however. He uses a sample of only 58 observations, considers return horizons only up to four years, and does not try to construct a data generating mechanism that replicates observed characteristics of the actual data.

These studies all agree that there are statistical pitfalls in evaluating long-run stock market performance. But it is striking how well the evidence for stock market

²³ The Monte Carlo results for the bottom part of Figure 1 are related to the results for the top part of Figure 1. If we had continuous data, so that the change in the dividend–price ratio to the next crossing of the mean was just minus the current demeaned dividend–price ratio, then the price regression coefficient for the bottom part and the dividend regression coefficient for the top part of Figure 1 would have to differ by one. In fact our data are not continuous but are measured annually, so the change in the dividend–price ratio to the next crossing of the mean exceeds the current demeaned dividend–price ratio in absolute value, and the two regression coefficients differ by slightly more than one. It is still true, however, that if the price regression coefficient is close to one then the dividend regression coefficient must be close to zero.

²⁴ Campbell and Shiller use a VAR that includes dividend growth, the dividend yield, and the ratio of smoothed earnings to prices. They construct a loglinearized approximation to the stock return from the dividend growth rate and the dividend yield.

predictability survives the various corrections and adjustments that have been proposed in this research.

V. Conclusion

We concluded in the 1998 version of this paper that the conventional valuation ratios, the dividend–price and price-smoothed–earnings ratios, have a special significance when compared with many other statistics that might be used to forecast stock prices. In 1998 these ratios were extraordinarily bearish for the US stock market. The ratios are even more so now.

These valuation ratios deserve a special place among forecasting variables because we have such a long time series of data on these ratios, and because they relate stock prices to careful evaluations of the fundamental value of corporations. Earnings have been calculated and reported by US corporations for over a hundred years for the express purpose of allowing us to judge intrinsic value. Dividend distribution decisions have been made by corporations for just as long with a sense that dividends should be set in such a way that they can reasonably be expected to continue.

Linear regressions of price changes and total returns on the log valuation ratios suggest substantial declines in real stock prices, and real stock returns below zero, over the next ten years. This result must of course be interpreted with caution. The valuation ratios are now so far from their historical averages that we have very little comparable historical data; our regressions extrapolate linearly from a relation between log valuation ratios and long-horizon returns that holds in historically normal times to get a prediction for the current, historically abnormal situation. It is quite possible that the true relation between log valuation ratios and long-horizon returns is nonlinear, in which case linear regression forecasts might be excessively bearish. But while this point may moderate the extreme pessimism of our linear regressions, it certainly does not support optimism about the stock market outlook.

It is also possible that forecasting relations that worked in the past will cease to work now. But these ratios are not forecasting variables that were discovered yesterday, *ex post*. They are *ex ante* forecasting relations that have been continually discussed over the last century.

The very fact that ratios have moved so far outside their historical range poses a challenge however, both to the traditional view that stock prices reflect rational expectations of future cash flows, and to our view that they are substantially driven by mean reversion. Observers of either persuasion must face the fact that something extremely unusual has occurred. In this situation a broad judgment of our position in history, of the uniqueness of recent technological advances and investment patterns, and of the state of market psychology assumes more than usual importance in judging the outlook for the stock market. There is no purely statistical method to resolve finally whether the data indicate that we have entered a new era, invalidating old relations, or whether we are still in a regime where ratios will revert to old levels. In our personal judgment, while we do not expect a complete return to traditional valuation levels, we still interpret the broad variety of evidence as suggesting a poor long-term outlook for the stock market.

References

- Bakshi, Gurdip S. and Zhiwu Chen, “Baby Boom, Population Aging, and Capital Markets,” *Journal of Business* 67: 165–202, April 1994.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, “A Model of Investor Sentiment,” *Journal of Financial Economics*, 1998.
- Barsky, Robert B. and J. Bradford De Long, “Why Does the Stock Market Fluctuate?” *Quarterly Journal of Economics* 107: 291–311, 1993.
- Blanchard, Olivier J., “Movements in the Equity Premium,” *Brookings Papers on Economic Activity* 2:75–118, 1993.
- Bond, Stephen R., and Jason G. Cummins, “The Stock Market and Investment in the New Economy,” *Brookings Papers on Economic Activity*, pp. 61–108, I 2000.
- Campbell, John Y., “Asset Prices, Consumption, and the Business Cycle”, in John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Vol. 1C, pp. 1231–1303, Amsterdam: Elsevier , 1999.
- Campbell, John Y. and John H. Cochrane, “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107: 205–251, April 1999.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ, 1997.
- Campbell, John Y., and Robert J. Shiller, “The Dividend–Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1:195–228, Fall 1988(a).
- Campbell, John Y., and Robert J. Shiller, “Stock Prices, Earnings, and Expected Dividends,” *Journal of Finance*, 43(3): 661–676, July 1988(b).
- Campbell, John Y. and Robert J. Shiller, “The Dividend–Ratio Model and Small Sample Bias: A Monte Carlo Study,” *Economics Letters*, 29: 325–331, 1989.
- Campbell, John Y. and Luis M. Viceira, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, forthcoming, Oxford University Press, 2001.
- Cole, Kevin, Jean Helwege, and David Laster, “Stock Market Valuation Indicators: Is This Time Different?” *Financial Analysts Journal*, 56–64, May/June 1996.
- Cowles, Alfred, and Associates, *Common Stock Indexes*, 2nd ed., Principia Press, Bloomington, 1939.
- Cummins, Jason G., “Comments and Discussion (on Hall paper),” *Brookings Papers on Economic Activity*, pp. 103–10, II 2000.
- Glassman, James K. and Kevin A. Hassett, *Dow 36, 000: The New Strategy for Profiting from the Coming Rise in the Stock Market*, Times Books, 1999.

- Goetzmann, William N., and Philippe Jorion, "Testing the Predictive Power of Dividend Yields," *Journal of Finance* 48:663–679, June 1993.
- Gordon, Robert J., "Interpreting the One Big Wave in U.S. Long-Term Productivity Growth," National Bureau of Economic Research Working Paper No. 7552, June 2000.
- Graham, Benjamin, and David L. Dodd, *Security Analysis*, First Edition, McGraw Hill, New York, 1934.
- Hall, Robert E., "E-Capital: The Link between the Stock Market and the Labor Market in the 1990s," *Brookings Papers on Economic Activity*, pp. 73–118, II 2000.
- Hall, Robert E., "Struggling to Understand the Stock Market," The 2001 Ely Lecture, American Economic Association, January 2, 2001, forthcoming, *American Economic Review*.
- Heaton, John and Deborah Lucas, "Stock Prices and Fundamentals," in Ben S. Bernanke and Julio J. Rotemberg, eds., *NBER Macroeconomics Annual*, 213–241, 1999.
- Kirby, Chris, "Measuring the Predictable Variation in Stock Returns," *Review of Financial Studies* 10: 579–630, Fall 1997.
- Liang, J. Nellie, and Steven A. Sharpe, "Share Repurchases and Employee Stock Options and Their Implications for S&P500 Share Retirements and Expected Returns," unpublished paper, Board of Governors of the Federal Reserve System, 1999.
- McGrattan, Ellen R., and Edward C. Prescott, "Is the Stock Market Overvalued?" National Bureau of Economic Research Working Paper Series #8077, January 2001, National Bureau of Economic Research, Cambridge MA.
- Modigliani, Franco and Richard A. Cohn, "Inflation, Rational Valuation, and the Market," *Financial Analysts Journal* 24–44, March–April 1979.
- Nelson, Charles R., and Myung J. Kim, "Predictable Stock Returns: The Role of Small Sample Bias," *Journal of Finance* 48: 641–661, June 1993.
- Nordhaus, William D., "Productivity Growth and the New Economy," Cowles Foundation Discussion Paper No. 1284, Cowles Foundation, Yale University, December 2000.
- Ritter, Jay R. and Richard S. Warr, "The Decline of Inflation and the Bull Market of 1982 to 1997," unpublished paper, University of Florida.
- Shiller, Robert J., "Human Behavior and the Efficiency of the Financial System," in John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Vol. 1C, pp. 1305–40, Amsterdam: Elsevier, 1999.
- Shiller, Robert J., *Irrational Exuberance*, Princeton University Press, Princeton New Jersey, 2000.
- Shiller, Robert J., *Market Volatility*, MIT Press, Cambridge, 1989.

Shiller, Robert J., “Price–Earnings Ratios as Forecasters of Returns: The Stock Market Outlook in 1996,” unpublished paper, Yale University, July 1996, <http://www.econ.yale.edu/~shiller>.

Siegel, Jeremy J., *Stocks for the Long Run*, 2nd ed., McGraw-Hill, 1998.

Vissing–Jørgenson, Annette, *Limited Stock Market Participation*, PhD thesis, MIT, 1998.

Figure 1. DIVIDEND GROWTH till the next time D/P crosses its mean

PRICE GROWTH till the next time D/P crosses its mean

Figure 2. 1-year DIVIDEND GROWTH vs D/P

1-year PRICE GROWTH vs D/P

Figure 3. 10-year DIVIDEND GROWTH vs D/P

10-year PRICE GROWTH vs D/P

Figure 4. S&P Composite Stock Data, January Values 1872–1997

Figure 5. 1-year GROWTH of 10-year MA(E) vs P/10-year MA(E)

1-year PRICE GROWTH vs P/10-year MA(E)

Figure 6. 10-year GROWTH of 10-year MA(E) vs P/10-year MA(E)

10-year PRICE GROWTH vs P/10-year MA(E)

Figure 7. Output/Hour and Real Earnings

10 Year Earnings and Output Growth

Figure 8. 10-year Productivity GROWTH vs P/10-year MA(E)

Figure 9. AUSTRALIA

CANADA

FRANCE

GERMANY

AUSTRALIA

CANADA

FRANCE

GERMANY

Figure 9 (continued): ITALY

JAPAN

NETHERLANDS

SPAIN

ITALY

JAPAN

NETHERLANDS

SPAIN

S W E D E N