Valuation Ratios and the Long-Run Stock Market Outlook

Ratios are extraordinarily bearish.

John Y. Campbell and Robert J. Shiller

JOHN Y. CAMPBELL is professor
of economics at Harvard University
(MA 02138).

ROBERT J. SHILLER is profes-
sor of economics at Yale University
(CT 06520).

When stock market valuation ratios are at extreme levels by historical standards, as dividend-price and price-earnings ratios are in the United States today, one naturally wonders what this means for the stock market outlook. It seems reasonable to believe that prices are not likely ever to drift too far from their normal relationships to indicators of fundamental value, such as dividends or earnings. Thus one might expect that when stock prices are very high relative to these indicators, as they are in 1997, prices will fall in the future to bring the ratios back to more normal historical levels.

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 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 on the basis of 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 efficient markets theory is that the stock market is 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 efficient markets theory is not to imply that

the dividend-price ratio will move beyond its historical range, or get stuck forever at the current extreme, it requires that the dividend-price ratio predicts future growth in dividends.[1]

WHAT DOES THE DIVIDEND-PRICE RATIO FORECAST?

Does the dividend-price ratio forecast future dividend movements as required by the efficient markets theory, or does it instead forecast future movements in stock prices? We address this question using a long-run annual U.S. data set that extends today's S&P 500 index back in time to 1872.[2]

The answer is given by the pair of scatterplots shown in Exhibit 1. The dividend-price ratio, measured as the previous year's dividend divided by the January stock price, is on the horizontal axis of both graphs. (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 4.73%.

In Panel A, 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 4.73%. More precisely, we measure the dividend growth rate from the year preceding the year shown until the year before the dividend-price ratio again crosses 4.73%. 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.[3]

The different years are indicated in the scatter diagram by two-digit numbers; an asterisk after a number denotes a nineteenth-century date. Since 1872, the dividend-price ratio has crossed its mean value twenty-nine times, with intervals between crossings ranging from one year to twenty years (the twenty-year interval is between 1955 and 1975). 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 1997. The implied forecast for dividend growth is the horizontal dashed line marked

EXHIBIT 1
DIVIDEND AND PRICE GROWTH

PANEL A. DIVIDEND GROWTH UNTIL NEXT TIME D/P CROSSES ITS MEAN VERSUS D/P
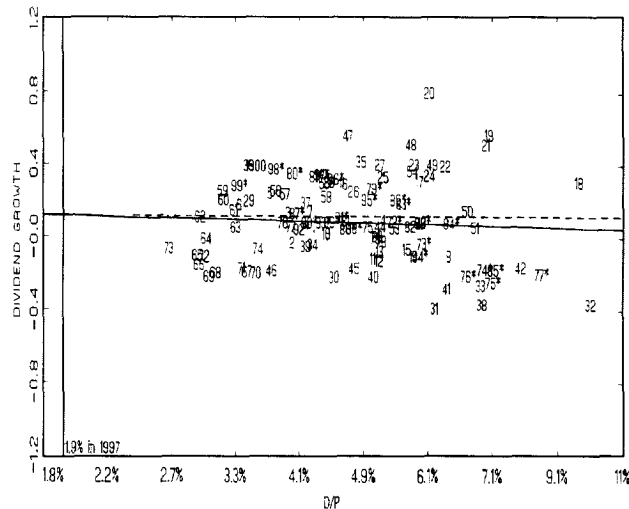
PANEL B. PRICE GROWTH UNTIL NEXT TIME D/P CROSSES ITS MEAN VERSUS D/P
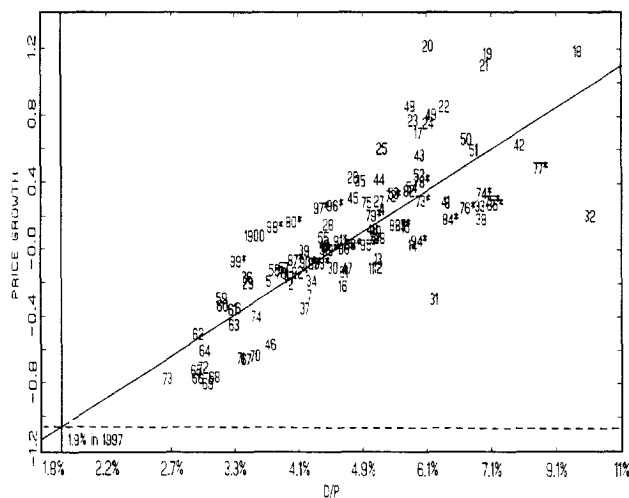
where the vertical line intersects the regression line.

It is obvious from Panel A 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 0.27%, indicating that only one-quarter of 1% 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 Panel B of Exhibit 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 crosses 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 64%.

We have answered the 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 1997, the dividend-price ratio was only 1.9%, well to the left of any points shown in Exhibit 1. (In the first six months of 1997, the stock market rose by about 20%, bringing the dividend-price ratio down even farther.) Panel B of Exhibit 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 1997 indicates that, the next time that the dividend-price ratio is back to its mean, the natural log real value of the stock market will be about 1.0 lower than it is today. Translated into percentage terms, this is equivalent to saying that the stock market will lose about two-thirds of its real value.

Can we take such a forecast seriously? What modifications should we make to such a forecast?

FIXED-HORIZON FORECASTS FROM THE DIVIDEND-PRICE RATIO

Exhibit 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 graphs 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; the historical range is from one

to twenty years.

We now show scatterplots like Exhibit 1, but where we change the vertical axis to show growth rates of dividends and prices over a fixed horizon. The horizon is one year in Exhibit 2, and ten years in Exhibit 3. We should expect to see a worse fit than in Exhibit 1, of course, since with these graphs we do not measure dividend and price growth rates over intervals when the ratio returns to its mean value.

Panel A of Exhibit 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 almost 15% 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.

Panel B of Exhibit 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 is visible in Exhibit 1.[4]

Many of the patterns of Exhibit 1 become apparent again in Exhibit 3, where the horizon is ten years rather than one year. Just as in Exhibit 1, there is only a very weak relation between the dividend-price ratio and subsequent ten-year dividend growth. In fact, the relation in Exhibit 3 is even less consistent with the efficient markets theory than the relation in Exhibit 1, because the Exhibit 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 Exhibit 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 15% for price growth.
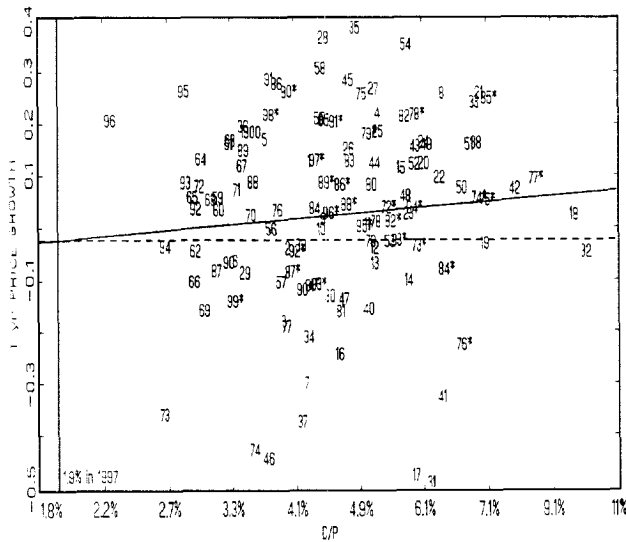
Given the low value for the dividend-price ratio at the start of 1997, the regression in Panel B of Exhibit

EXHIBIT 2
ONE-YEAR GROWTH MEASURES

PANEL A. ONE-YEAR DIVIDEND GROWTH VERSUS D/P

PANEL B. ONE-YEAR PRICE GROWTH VERSUS D/P

EXHIBIT 3
TEN-YEAR GROWTH MEASURES

PANEL A. TEN-YEAR DIVIDEND GROWTH VERSUS D/P
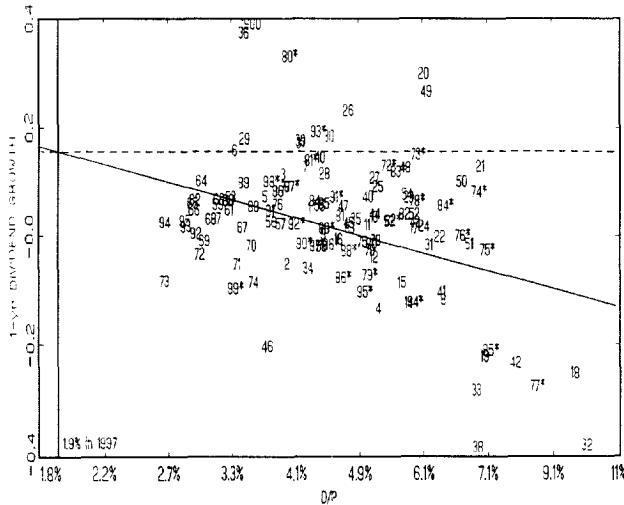
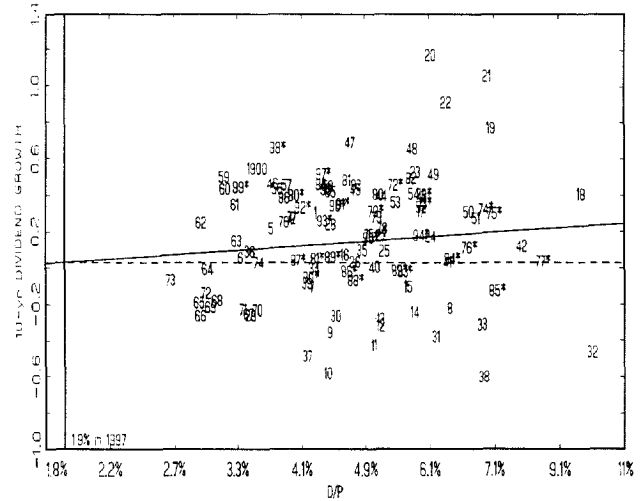**PANEL B. TEN-YEAR PRICE GROWTH VERSUS D/P**

3 implies a change of −0.475 in the log real stock price over the next ten years. This corresponds to a 38% loss of real value.[5]
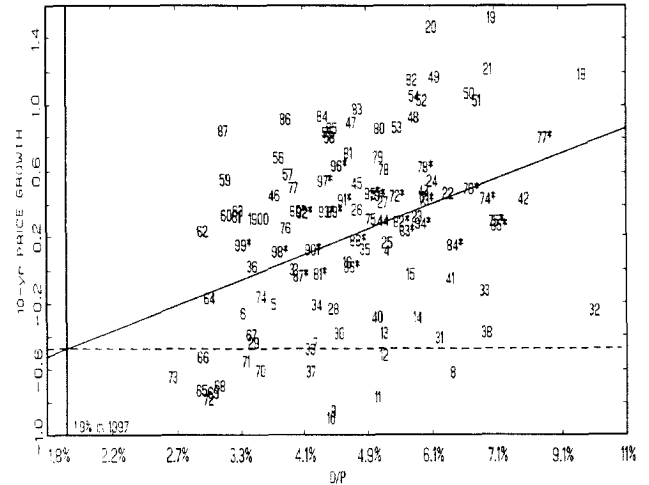
The poorer fit in Exhibits 2 and 3, when the vertical axis shows a fixed horizon, reflects the fact that we do not know when the dividend-price ratio will be restored to conventional levels. Part of the reason for
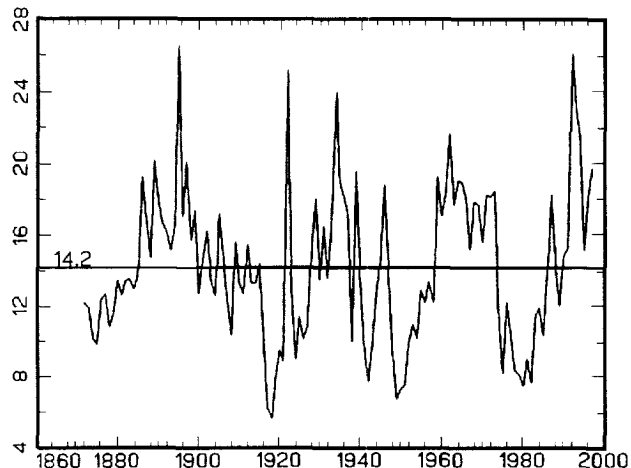
this may be that the numerator of the ratio, the dividend, is just not the most accurate available measure of fundamental value. We must consider ratios constructed from some other measures of fundamental value.
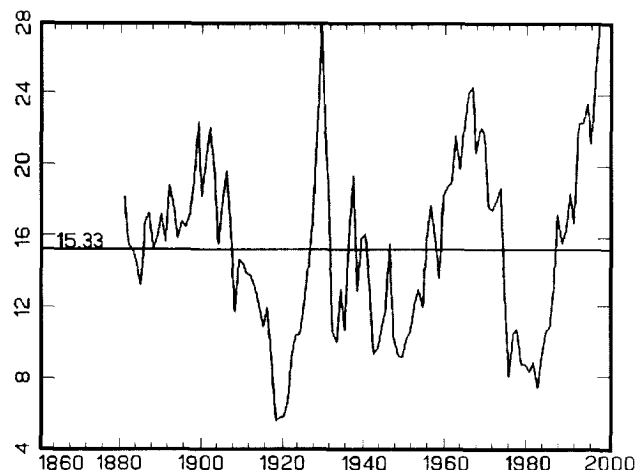
ALTERNATIVE VALUATION RATIOS

Exhibit 4 illustrates some key valuation ratios in

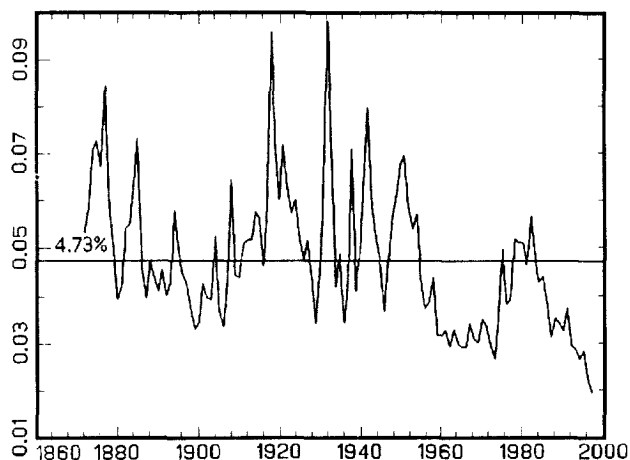**EXHIBIT 4**
S&P COMPOSITE STOCK DATA JANUARY VALUES 1872-1997
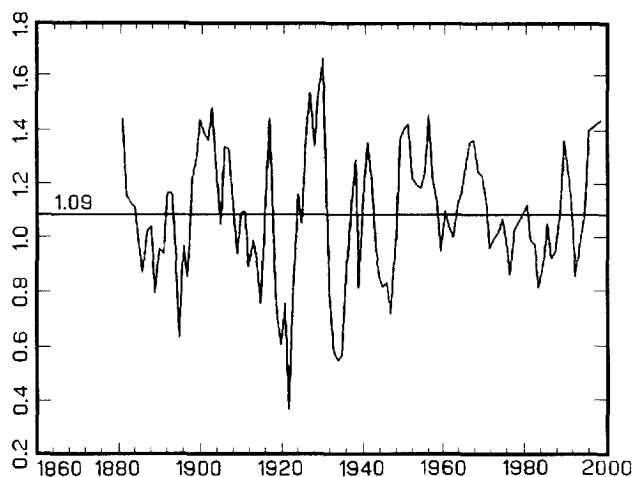
PANEL A. P/E

PANEL C. P/10-YEAR MA(E)

PANEL B. D/P

PANEL D. E/10-YEAR MA(E)

our long-run annual U.S. data set. Panel A shows the price-earnings ratio, calculated using the January stock price in each year divided by the level of earnings from the previous year. Panel B 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.

Exhibit 4 illustrates the fact that price-earnings ratios have normally moved in a range from 8 to about

20, with a mean of 14.2 and occasional spikes down as far as 6 or up as high as 26. At the beginning of 1997, the price-earnings ratio was high, at 20, but not at a record level. Dividend-price ratios have normally moved in a range from 3% to about 7%, with a mean of 4.73% and occasional movements up to almost 10%. Very recently the dividend-price ratio has fallen to a record low below 2%.

Since stock price increases drive up price-earnings ratios and drive down dividend-price ratios, it is not surprising that the two series in Exhibit 4 general-

ly 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 famous 1934 textbook, Security Analysis, say 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" [1934, p. 452]. Following their advice, we smooth earnings by taking an average of real earnings over the past ten years.[6]

Exhibit 4, Panels C and D, shows the ratio of the January real stock price to smoothed real earnings from the previous year and the ratio of current real earnings to smoothed real earnings. The price-smoothed earnings ratio in Panel C 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 15.3, but the record high of 28 now appears at the start of 1997. The only previous year with a comparably high ratio is 1929.

## FORECASTS FROM THE PRICE-SMOOTHED EARNINGS RATIO

Exhibits 5 and 6 have the same format as Exhibits 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 less than 4% over one year and over ten years. The ratio is a good forecaster of ten-year growth in stock prices, with an R^2 statistic of 37%. The fit of this relation is substantially better than we found for the dividend-price ratio in Exhibit 3.[7]

Noting that the price-smoothed earnings ratio for January 1997 is a record 28, the regression illustrated in Exhibit 6 is predicting a decline of 0.5 in the log real stock price. In percentage terms, it is predicting that the real value of the market will be 40% lower in

EXHIBIT 5
SMOOTHED ONE-YEAR GROWTH MEASURES

PANEL A. ONE-YEAR GROWTH OF TEN-YEAR MA(E) VERSUS P/TEN-YEAR MA(E)
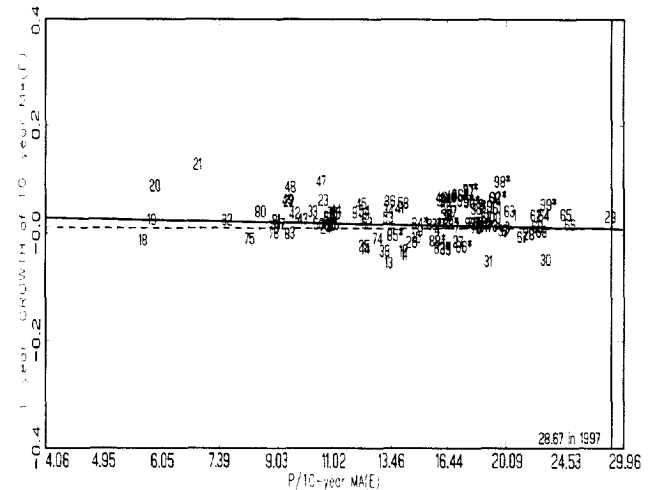
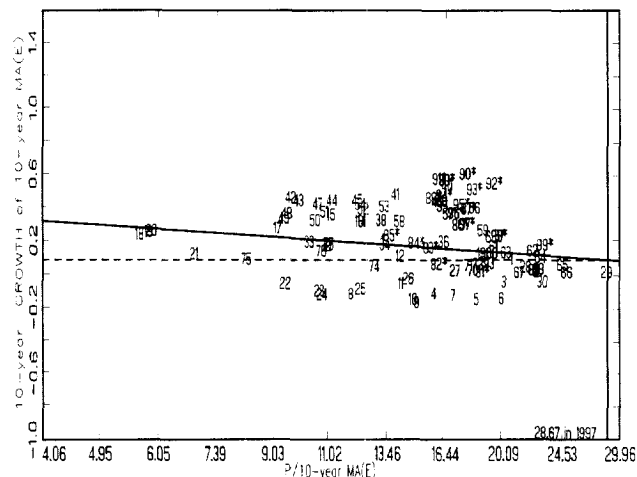### PANEL B. ONE-YEAR PRICE GROWTH VERSUS P/TEN-YEAR MA(E)
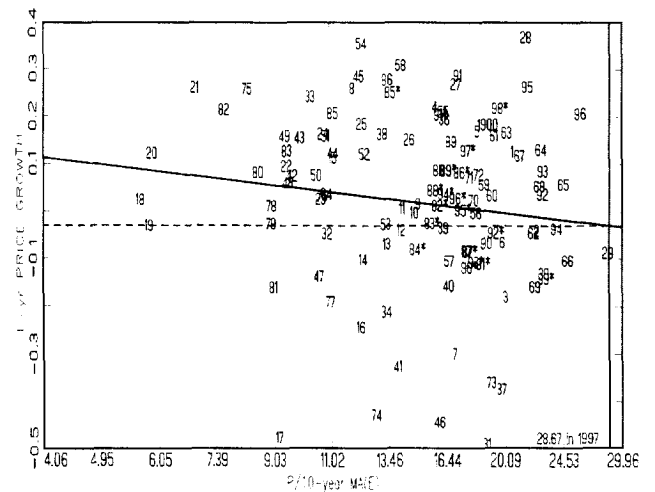
ten years than it is today. The corresponding forecast for the cumulative continuously compounded real stock return is −15% over ten years.

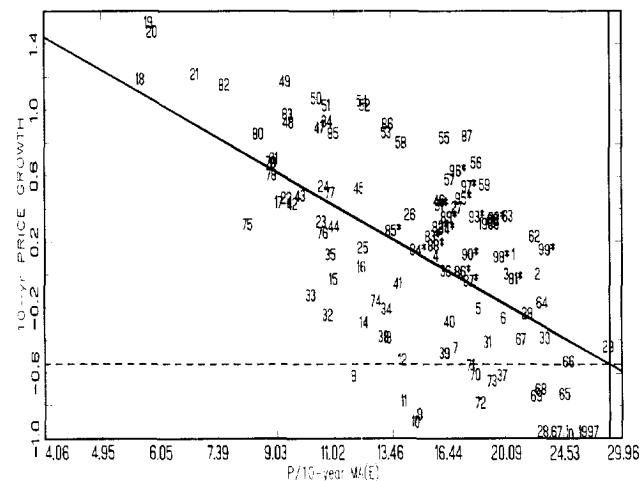## WILL VALUATION RATIOS REMAIN IN THEIR HISTORICAL RANGES?

Over the past century, the American economy has been transformed in many fundamental ways. Agriculture gave way to industry, and industry has

EXHIBIT 6
SMOOTHED TEN-YEAR GROWTH MEASURES

PANEL A. TEN-YEAR GROWTH OF TEN-YEAR MA(E) VERSUS P/TEN-YEAR MA(E)

PANEL B. TEN-YEAR PRICE GROWTH VERSUS P/TEN-YEAR MA(E)

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.

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 twentieth 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 U.S. economy.

FINANCIAL MARKET CHANGES

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. Companies can repurchase their stock, for instance, as a tax-favored alternative to paying dividends. Repurchases transfer cash to the 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 the record low dividend-price ratio in 1997.

One way to correct 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] do this for S&P 500 firms over the period

1975-1995, and find that dividend-price ratios should be adjusted upward significantly during the mid-1980s and the mid-1990s. Their adjustment increases the 1996 dividend-price ratio by 80 basis points, from 2.2% to almost 3.0%. Shulman, Brown, and Narayanan [1997] calculate a similar adjustment for 1997.

A glance at Exhibit 4 shows that an adjustment of this magnitude brings the dividend-price ratio back 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 1997 valuation ratios.

Shulman, Brown, and Narayanan [1997] make a further important point. The adjustments just described assume 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 are issuing shares below market value as part of their employee stock option incentive plans. Thus they may issue more shares than they repurchase, even if the net cost of repurchases and issues is positive. Shulman, Brown, and Narayanan argue that this is in fact the case for S&P 500 firms in 1997, so that the dividend-price ratio should be adjusted downward rather than upward.

The Baby Boom and the Demand for Stock

Many observers also argue that there has been a secular shift in the attitudes of the investing public toward 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 twenty-first century — and may even move farther outside their historical ranges if the demographic effect strengthens.

In support of this argument, it has been pointed out that the dividend-price ratio shows some evidence of trend decline during the entire period since World War II. The appearance of long-run stability in this ratio in Exhibit 4 would be much weaker if the figure began in the mid-twentieth century rather than in 1872.[8]

While there may be some validity to this argument, it does not necessarily contradict the pessimistic stock market outlook we presented earlier. The demographic 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 here, 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 Exhibit 4; high price-earnings and price-smoothed earnings ratios and low dividend-price ratios are characteristic of periods when real earnings have been growing rapidly (such as the 1920s, 1960s, and mid-1990s), 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.[9]

Inflation

Other observers have argued that today's high stock prices can be justified by the steady decline in inflation that has taken place during the 1990s. 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.[10] Thus it should not be surprising to see high stock prices, given the low inflation of the mid-1990s.

There are two weaknesses in this argument. First, there has been no clear downward trend in inflation or

long-term bond yields since the beginning of 1994. Thus it is hard to explain the recent rise in the stock market by any large change in the outlook for inflation.

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. 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.

International Evidence

We have emphasized that in the U.S. 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 U.S. experience anomalous?

Unfortunately, very little long-term information is 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 Exhibit 4 that since the early 1970s the time series plot of the U.S. 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. Exhibit 7 presents scatterplots like Exhibits 2 and 3, but with quarterly data and a four-year horizon. The dividend-price ratio appears on the horizontal axis of each 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.[11]

The countries in Exhibit 7 fall into two main groups. The English-speaking countries, Australia, Canada, and the U.K., along with some of the smaller continental European countries, the Netherlands, Spain, Sweden, and Switzerland, behave over this short sample period very much like the U.S. The dividend-price ratio is positively associated with subsequent price growth, and shows little relation to subsequent dividend growth.

The large continental European countries, France, Germany, and Italy, show a very different pattern over this sample period. In these countries a high dividend-price ratio is associated with weak subsequent dividend growth, just as the efficient markets theory would imply. There is little relation between the dividend-price ratio and subsequent price growth.

Japan is an intermediate case. The dividend-price ratio appears to have been associated with both subsequent dividend growth and subsequent price growth.

The weight of the evidence from these other countries is limited, given the short sample period. Note, for example, that in both France and Germany we would have observed a substantial positive relation between the dividend-price ratio and subsequent price growth if the dividend-price ratio had been lower in these countries in the early 1970s, closer to the U.S. level. The difference between the U.S. and the French or German scatterplots can be traced to a single business downturn around the time of the first oil crisis. No one country shown in Exhibit 7 can provide very substantial evidence about long-run stock market behavior.

Nonetheless, since most of our twelve countries show the same pattern as the long-term U.S. data, these results seem generally consistent with the importance of valuation ratios. Note also that in most of these countries the dividend-price ratio is now very low, and the fitted regression lines predict substantial four-year declines in stock prices.
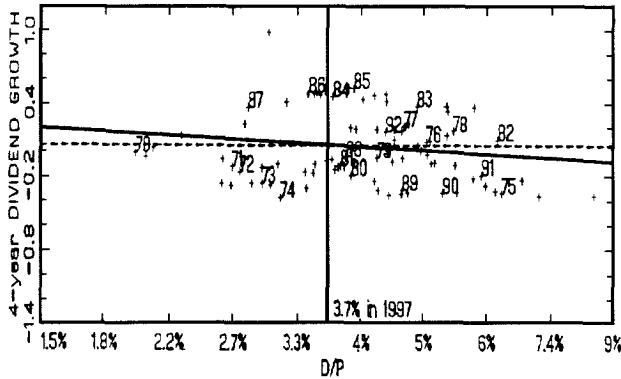
SOME STATISTICAL PITFALLS

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 Exhibit 7), 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 ran-
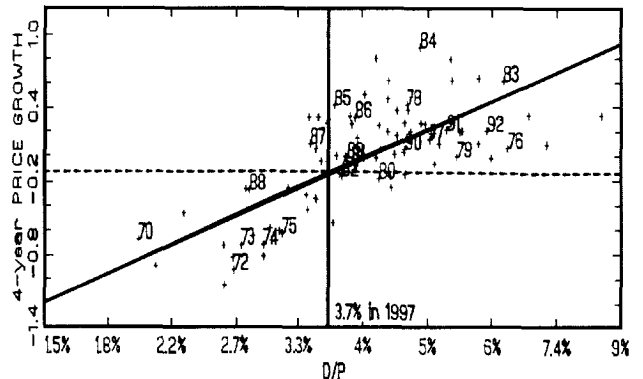
EXHIBIT 7
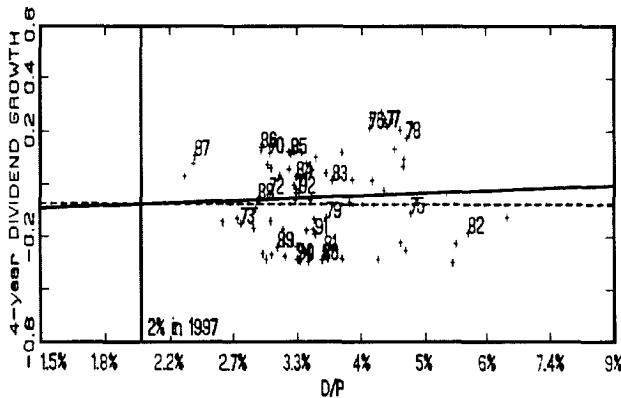## INTERNATIONAL PATTERNS IN FOUR-YEAR DIVIDEND GROWTH AND PRICE GROWTH
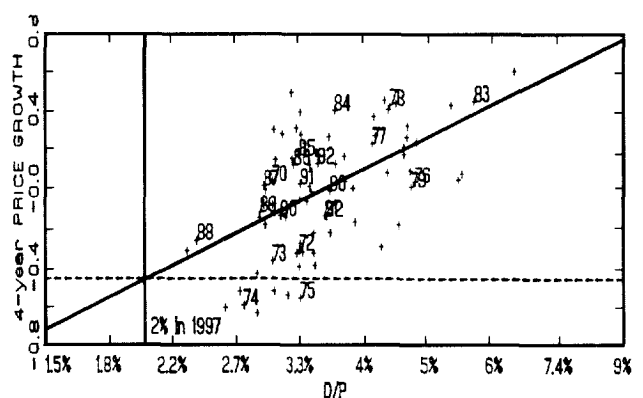
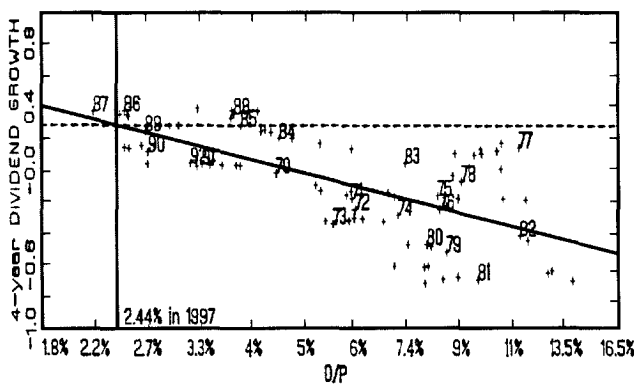**AUSTRALIA — DIVIDEND GROWTH VERSUS D/P**   **PRICE GROWTH VERSUS D/P**

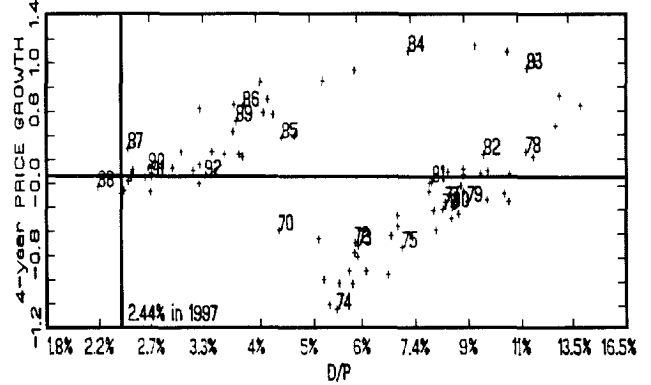**CANADA — DIVIDEND GROWTH VERSUS D/P**   **PRICE GROWTH VERSUS D/P**

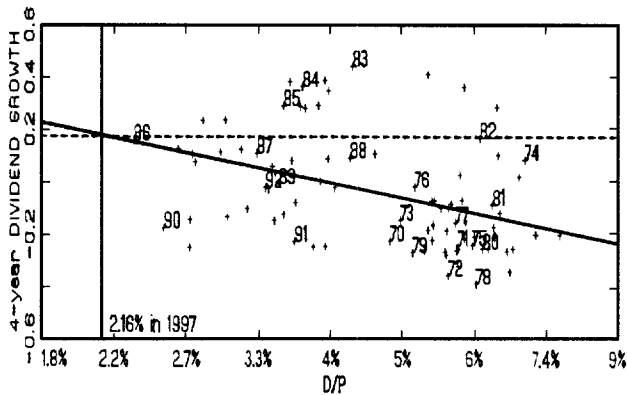**FRANCE — DIVIDEND GROWTH VERSUS D/P**   **PRICE GROWTH VERSUS D/P**

dom 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 Exhibit 1. We noted that the slope of the regression line in Panel A of the figure, predicting log real dividend growth over the time interval to the next crossing of the mean of the dividend-price ratio, is 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?
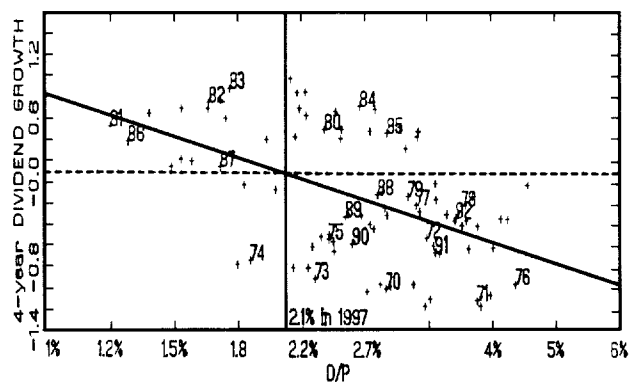
EXHIBIT 7
CONTINUED

GERMANY — DIVIDEND GROWTH VERSUS D/P

PRICE GROWTH VERSUS D/P
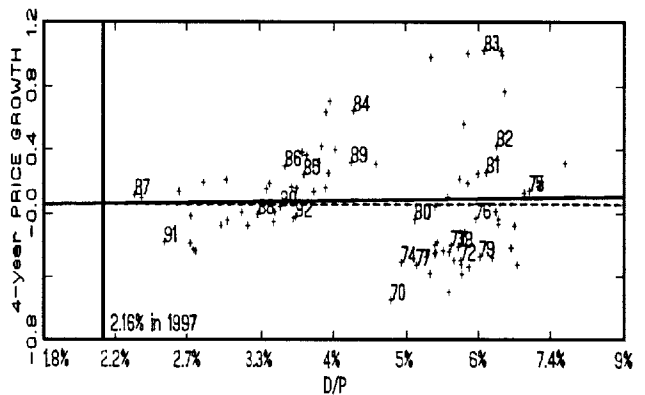
ITALY — DIVIDEND GROWTH VERSUS D/P

PRICE GROWTH VERSUS D/P

JAPAN — DIVIDEND GROWTH VERSUS D/P
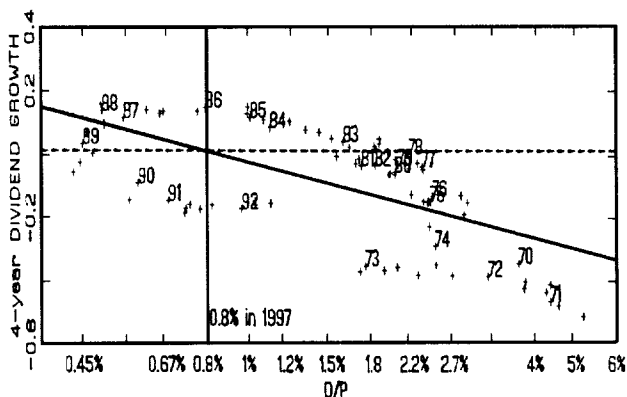
PRICE GROWTH VERSUS D/P
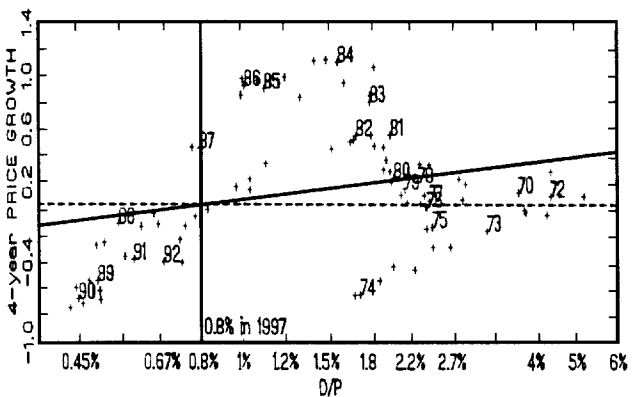
We perform a simple Monte Carlo experiment to study this issue. We construct artificial data in which the dividend-price ratio does not forecast future price changes over any fixed horizon. In other words, we generate data that satisfy the efficient markets prediction that the real stock price is a random walk.[12] Also,

we generate the data to match several important characteristics of the actual annual U.S. data.

We begin 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 correct the regression coefficient for small-

EXHIBIT 7
CONTINUED

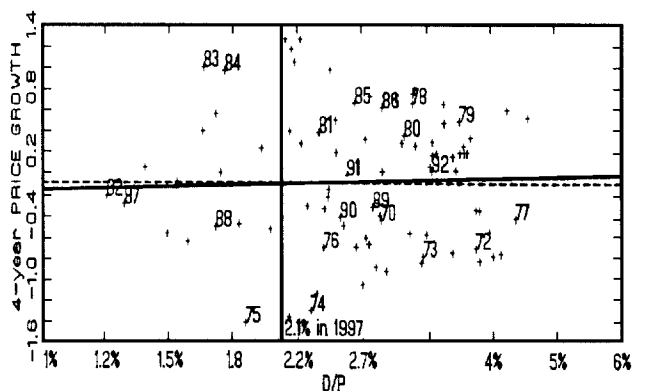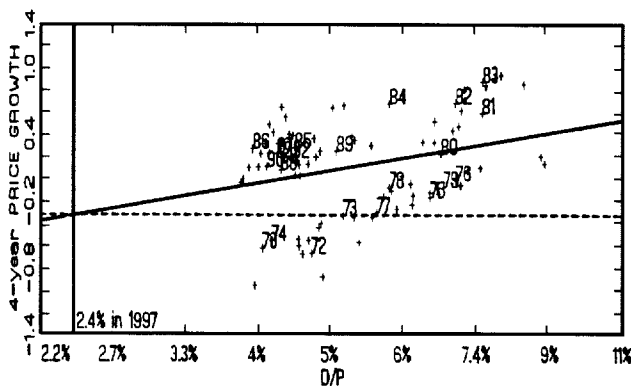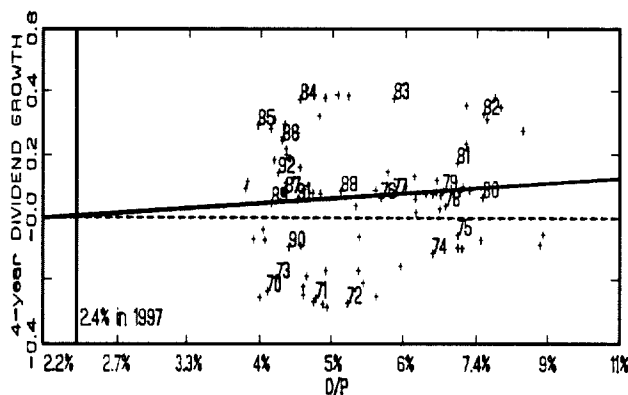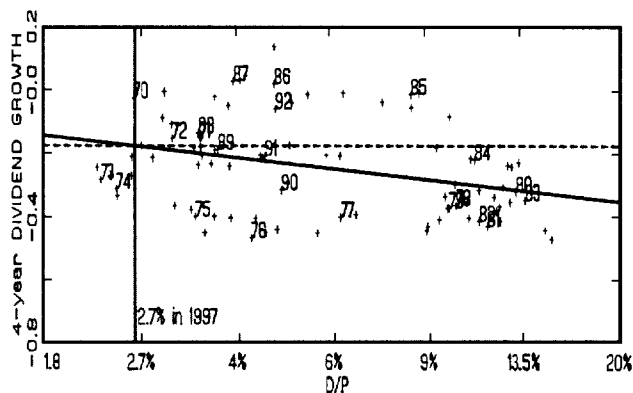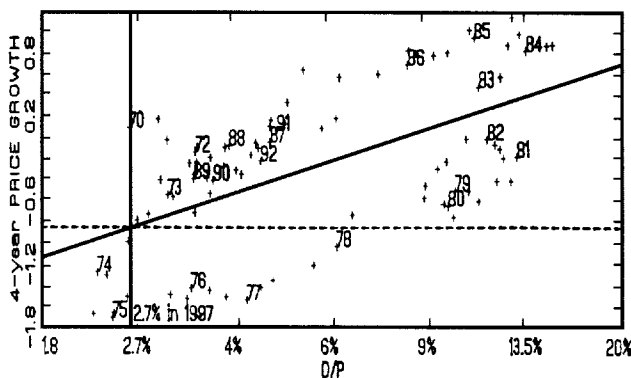THE NETHERLANDS — DIVIDEND GROWTH VERSUS D/P   PRICE GROWTH VERSUS D/P

SPAIN — DIVIDEND GROWTH VERSUS D/P   PRICE GROWTH VERSUS D/P

SWEDEN — DIVIDEND GROWTH VERSUS D/P   PRICE GROWTH VERSUS D/P
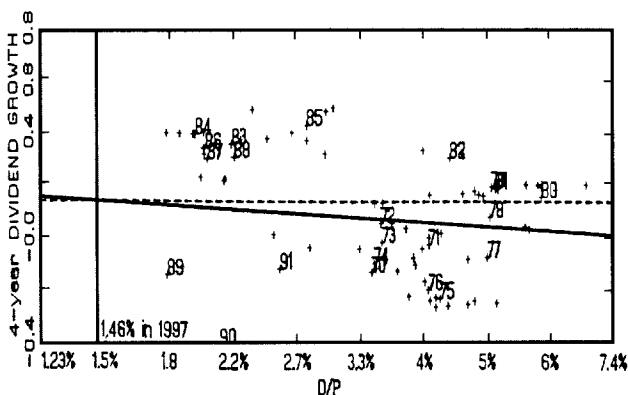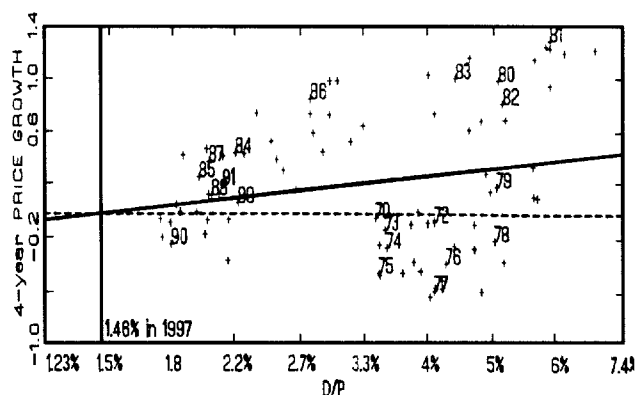
sample bias using the Kendall correction, obtaining a coefficient of 0.81. Using a random normal number generator with the estimated standard error of the error term in the bias-corrected regression, and using a random normal starting value whose variance equals the unconditional variance for this AR(1) model, we gen-

erate 125 observations of a simulated AR(1) log dividend-price ratio.

Next, we generate 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

EXHIBIT 7
CONTINUED

SWITZERLAND — DIVIDEND GROWTH VERSUS D/P
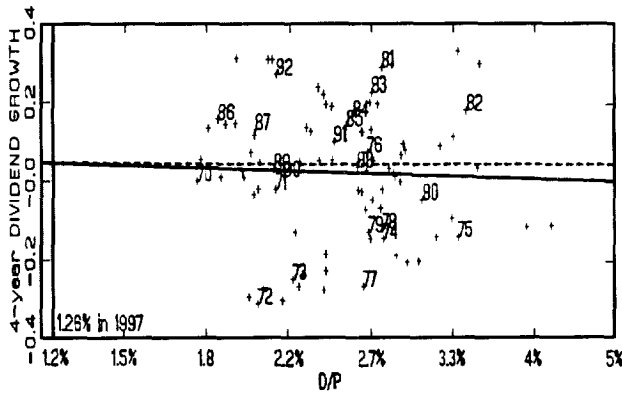
PRICE GROWTH VERSUS D/P
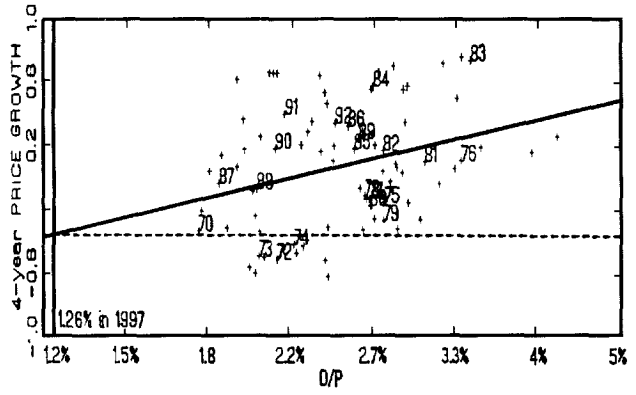
U.K. — DIVIDEND GROWTH VERSUS D/P
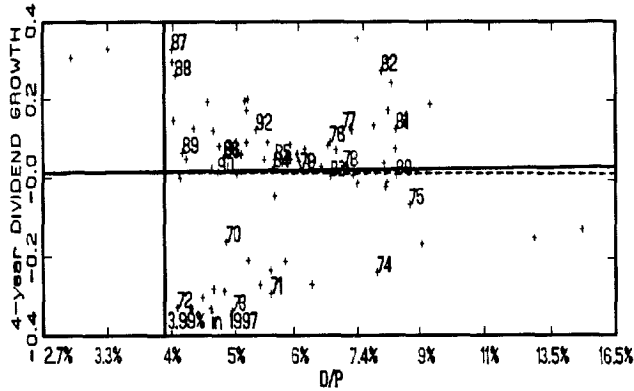
PRICE GROWTH VERSUS D/P
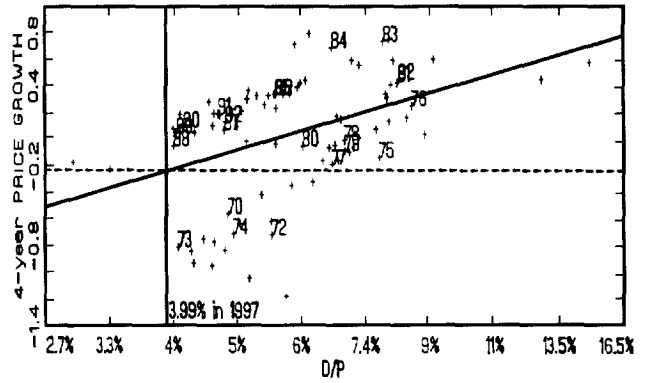
U.S. — DIVIDEND GROWTH VERSUS D/P
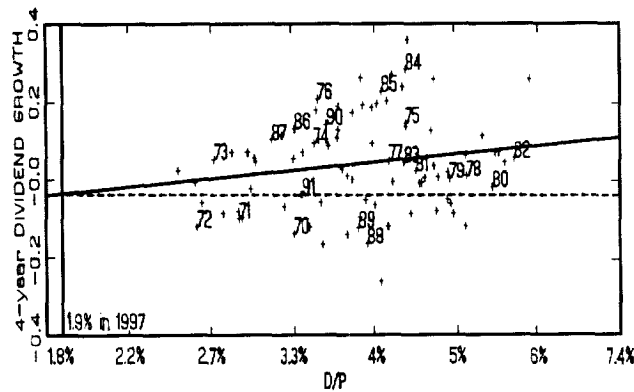
PRICE GROWTH VERSUS D/P
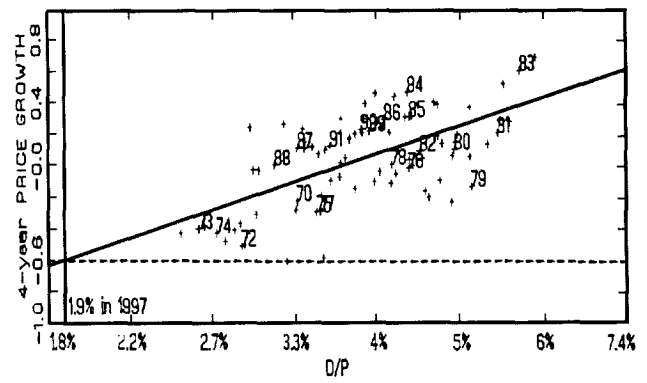
price. In the actual data, changes in the stock price and in the dividend-price ratio have a negative covariance; we also match this covariance in our artificial data.

Finally, we generate a log real previous-year dividend by adding the log dividend-price ratio and the log stock price.

We repeat this exercise 100,000 times. In each iteration, we use the artificial data to produce scatters and regression lines based on 125 observations like those shown in Panel A of Exhibit 1. We find that the average number of crossings of the mean of the dividend-price ratio is 26.5, not far from the 29 observed

with our actual data. But in 100,000 iterations, we find that the slope of the regression line shown in Panel A of Exhibit 1 is almost always much more negative than the estimated slope with the actual data. The estimated slope with the artificial data is greater than the estimated slope with actual data (−0.04) only 0.02%, two-hundredths of 1%, of the time.

The estimated regression coefficient in these Monte Carlo iterations tends 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 Panel A of Exhibit 1 is indeed anomalous from the standpoint of efficient markets theory.

Next we use the change in the log real stock price as the dependent variable in the Monte Carlo experiment, so that in each iteration we estimate the regression line shown in Panel B of Exhibit 1. In 100,000 iterations, we never once obtain a regression coefficient as great as the slope coefficient of 1.25 shown in Exhibit 1. While the average estimated slope coefficient in the Monte Carlo experiments is positive, the average value is only 0.18, far below the estimated coefficient with actual data.[13]

Other Monte Carlo experiments relevant to judging the results in this article 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.[14]

Nelson and Kim find that ten-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 non-stationary 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 anomalous.

Kirby [1997] uses Monte Carlo methods to further illustrate biases that can arise in conventional statistical tests of market efficiency. His results are not very relevant to our regressions, however. He uses a sample of only fifty-eight 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. Yet it is striking how well the evidence for stock market predictability survives the various corrections and adjustments that have been proposed in this research.

CONCLUSION

We think 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. Today these ratios are extraordinarily bearish for the U.S. stock market.

These valuation ratios deserve a special place among forecasting variables because we have such a long time series of data on them, and because they relate stock prices to careful evaluations of the fundamental value of corporations. Earnings have been calculated and reported by U.S. 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 valuation ratios suggest substantial declines in real stock prices, and real stock returns close to 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 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 valuation ratios and long-horizon returns is non-linear, 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, and they have been examined continually over the last century.

There may be special circumstances now that will change the historical relations between the valuation ratios and subsequent stock market performance. But there have always been special circumstances, circumstances that are adduced every time the ratios have been at extremes and that have in the past allowed people to fail to heed the message of the ratios.

ENDNOTES

This article is based on joint testimony before the Board of Governors of the Federal Reserve System, December 3, 1996, and on material in Shiller [1996]. The authors acknowledge the able research assistance of Elena Ranguelova and the helpful comments of Paul Samuelson.

[1]We are somewhat oversimplifying the efficient markets theory here. First, we are assuming that the equilibrium rate of return required by investors is constant; many recent versions of the efficient markets theory allow this required rate of return to vary over time.

Second, even with this assumption, 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 efficient markets 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. This effect is small, however, and in practice forecasts of returns and forecasts of price changes are very similar.

See Campbell and Shiller [1988a] for a careful analysis of dividend forecasts within the context of a log-linearized mathematical representation of the efficient markets theory. See also

Campbell, Lo, and MacKinlay [1997, Chapter 7] for a recent textbook exposition.

[2]The data in this article use the January Standard & Poor's 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.

[3]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, twenty-nine non-overlapping time intervals in our sample, so the data are not insubstantial. Statistical tests of the significance of analogous relations with fixed horizons, taking account of the overlapping intervals, are reported in Campbell and Shiller [1988b, 1989].

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

[5]As we have noted, stock returns differ from stock price changes because they include the direct contribution of dividends. Exhibit 3 implies an unusually poor 1997 outlook for stock returns, for three reasons. First, dividends are initially low relative to prices. Second, Panel A of Exhibit 3 shows that dividends are predicted to grow slowly over the next ten years. Third, Panel B of Exhibit 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 Panel B of Exhibit 3, but with a better fit (an R^2 statistic of almost 24% rather than 15%). The cumulative continuously compounded ten-year return forecast implied by the January 1997 dividend-price ratio is −16%.

[6]We first look at smoothed real earnings in Campbell and Shiller [1988b]. There we average log real earnings rather than levels of real earnings (that is, we use a geometric rather than an arithmetic average), but this makes little difference to the results. We also compare ten-year and thirty-year moving averages of earnings, and find that they have similar properties.

[7]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 Exhibit 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 46%, while a regression of ten-year returns on the dividend-price ratio has an R^2 statistic of 24%.

[8]Blanchard [1993] emphasizes the post-war decline trends 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.

[9]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 [1997] have argued 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 [1997] argue 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 its absolute level.

[10]A chart illustrating this common movement can be found in "Investment Strategy Chartbook" [1996, p. 16].

[11]See Campbell [1988] for a more detailed analysis of these international data.

[12]As before, we are oversimplifying the efficient markets theory by ignoring the distinction between price changes and returns. In Campbell and Shiller [1989] we generate 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.

[13]The Monte Carlo results for Panel B of Exhibit 1 are related to the results for Panel A. If we had continuous data, so that the change in the dividend-price ratio to the next crossing of the mean is just minus the current de-meaned dividend-price ratio, then the price regression coefficient for Panel B and the dividend regression coefficient for Panel A 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 de-meaned 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.

[14]Campbell and Shiller use a VAR that includes dividend growth, the dividend yield, and the ratio of smoothed earnings to prices. They construct a log-linearized approximation to the stock return from the dividend growth rate and the dividend yield.

REFERENCES

Bakshi, Gurdip S., and Zhiwu Chen. "Baby Boom. Population Aging, and Capital Markets." *Journal of Business*, 67 (April 1994). pp. 165-202.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny. "A Model of Investor Sentiment." Working paper, University of Chicago. 1997.

Barsky, Robert B., and J. Bradford De Long. "Why Does the Stock Market Fluctuate?" *Quarterly Journal of Economics*, 107 (1993). pp. 291-311.

Blanchard. Olivier J. "Movements in the Equity Premium." *Brookings Papers on Economic Activity*. 2 (1993), pp. 75-118.

Campbell, John Y. "Asset Prices, Consumption, and the Business Cycle." In John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*. Amsterdam: North-Holland, 1998.

Campbell, John Y., and John H. Cochrane. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." Working paper, Harvard University, April 1997.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay. *The Econometrics of Financial Markets*. Princeton: Princeton University Press, 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 (Fall 1988a), pp. 195-228.

——. "The Dividend-Ratio Model and Small Sample Bias: A Monte Carlo Study." *Economics Letters*, 29 (1989), pp. 325-331.

——. "Stock Prices, Earnings, and Expected Dividends." *Journal of Finance*, 43, 3 (July 1988b), pp. 661-676.

Cole. Kevin, Jean Helwege, and David Laster. "Stock Market Valuation Indicators: Is This Time Different?" *Financial Analysts Journal*. May/June 1996, pp. 56-64.

Cowles, Alfred, and Associates. *Common Stock Indexes*, 2nd ed. Bloomington: Principia Press, 1939.

Goetzmann, William N., and Philippe Jorion. "Testing the Predictive Power of Dividend Yields." *Journal of Finance*, 48 (June 1993), pp. 663-679.

Graham, Benjamin, and David L. Dodd. *Security Analysis*, first edition. New York: McGraw-Hill, 1934.

"Investment Strategy Chartbook." Goldman Sachs, November 1996.

Kirby. Chris. "Measuring the Predictable Variation in Stock Returns." *Review of Financial Studies*, 10 (Fall 1997), pp. 579-630.

Modigliani. Franco, and Richard A. Cohn. "Inflation, Rational Valuation, and the Market." *Financial Analysts Journal*, March-April 1979, pp. 24-44.

Nelson. Charles R., and Myung J. Kim. "Predictable Stock Returns: The Role of Small Sample Bias." *Journal of Finance*, 48 (June 1993), pp. 641-661.

Shiller, Robert J. *Market Volatility*. Cambridge: MIT Press, 1989.

——. "Price-Earnings Ratios as Forecasters of Returns: The Stock Market Outlook in 1996." Working paper, Yale University, July 1996.

Shulman, David, Jeffrey Brown, and Nari Narayanan. "Share Repurchases: Less Than Meets the Eye." Salomon Brothers, New York. May 29, 1997.