betting against the market during this recent bubble. They seem to be going right along with it. One of the biggest arguments for market efficiency has been that if the market is inefficient, why are the smart people still investing in the market. So, the question of how expert opinion can be biased will be one of the focal points of this talk.

The Recent Market Bubble

Figure 1 is the Nasdaq Composite Index in real terms from October 1984 to October 2001. Anyone who is thinking about the equity premium puzzle ought to reflect on what an event like the recent bubble we have had implies about the models of human rationality that underlie the equity premium puzzle. There has never been a more beautiful picture of a speculative bubble and its burst than in the Figure 1 chart of the Nasdaq; the price increase appears to continue at an ever increasing rate until March 2000; then, there is a sudden and catastrophic break, and the index loses a great deal of its value. We will have to reflect on what could have driven such an event before we can be comfortable with the economic models that imply a high degree of investor consistency and rationality.

Figure 2 shows the same speculative bubble from 1999 to late 2000 in the monthly real price and earnings of the S&P Composite Index since 1871. This bubble is almost unique; the only other one like it for the S&P Composite occurred in the 1920s; we could perhaps add the period just before the mid-1970s as a similar event. So, because we have a record of only two (possibly three) such episodes in history, a lot of short-run historical analysis may be misleading. We are in very unusual times, and this circumstance is obvious when we look at Figure 2.

The bubble that was seen in the late 1990s was not entirely confined to the stock market. Real estate prices also went up rapidly then. Karl Case[2] and I have devised what we call the "Case–Shiller Home Price Indexes" for many cities in the United States. **Figure 3** is our Los Angeles index on a quarterly basis from the fourth quarter of 1975 to the second quarter of 2001. (The smoothness in price change is not an artifact; real estate price movements tend to be smooth through time. The real estate market is different from the stock market.) Figure 3 tells an interesting and amazingly simple story. The two recessions over the period—1981–1982 and 1990–1991—are easy to see. Los Angeles single-family home prices were trending up when the 1981–82 recession hit. Then, although nominal home prices did not go down, prices did drop in real terms. After that recession, prices moved up again, only to fall again in the 1990–91 recession. Following that recession, prices soared back up. In the fall of 2001, we are again entering a recession. So, our prediction is that home

[2] Of Wellesley College, Massachusetts, and the real estate research firm of Case Shiller Weiss, Inc.

Figure 1. Real Nasdaq Composite, October 1984–October 2001

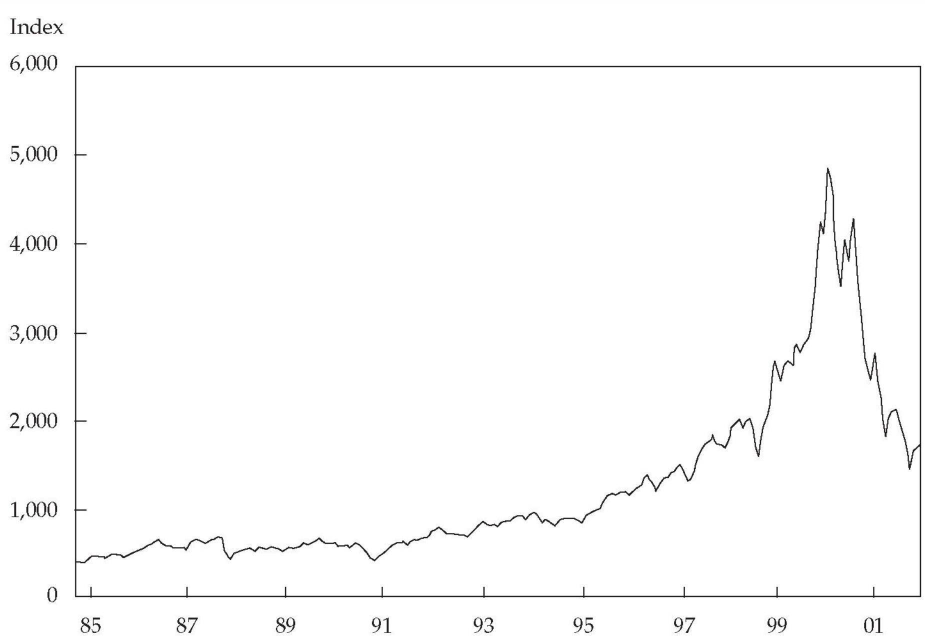

Figure 2. S&P Composite: Real Price and Earnings, January 1871–2001

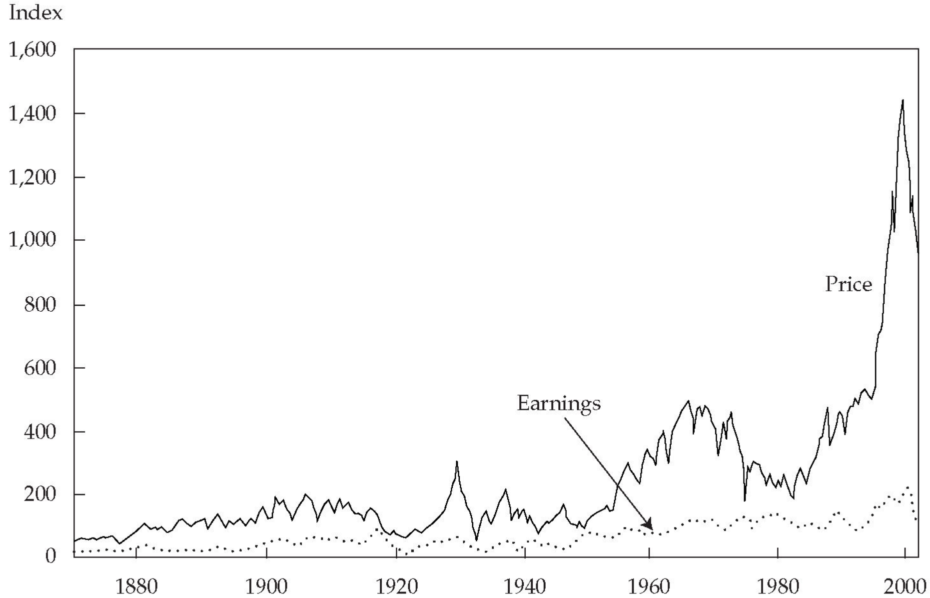

Note: Measured monthly.

Figure 3. Case–Shiller Home Price Index: Los Angeles Single-Family Home Prices, Fourth Quarter 1975 to Second Quarter 2001

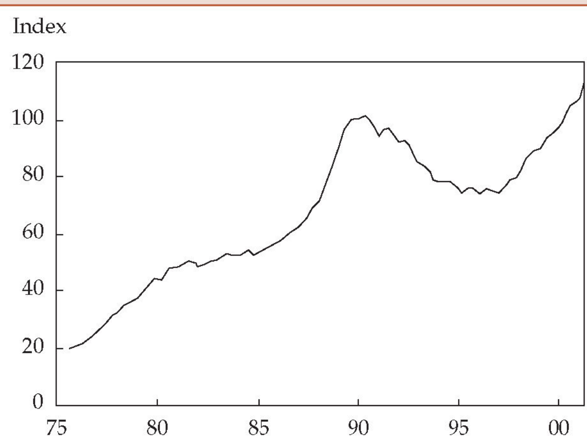

Note: Measured quarterly.

Figure 4. P/E for the S&P Composite, January 1881–October 2001

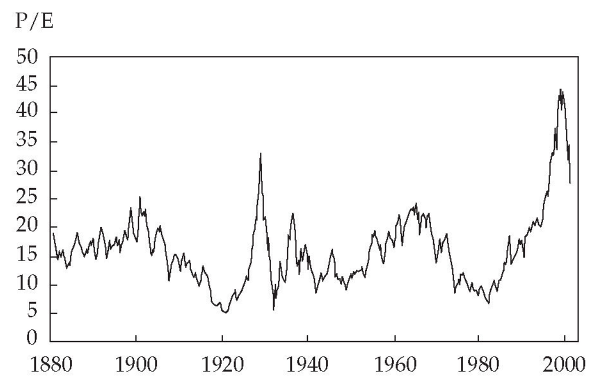

Note: P/E calculated as price over 10-year lagging earnings (a calculation recommended by Graham and Dodd in 1934).

prices may trend lower as a result. We do not expect to see in the market for homes a sharp bubble and burst pattern such as we saw in the Nasdaq, but we might well see some substantial price declines.

Figure 4, the S&P Composite P/E for 1881 to 2001, shows once again the dramatic behavior in the stock market recently, behavior matched only by the market of the late 1920s and (to a lesser extent) around 1900 and the 1960s.

Figure 5 is a scatter diagram, which John Campbell and I devised, depicting the historical negative correlation between P/Es and subsequent 10-year returns. Figure 5 shows how the S&P Composite P/E predicts future S&P Composite returns. The P/E is now around the 1929 level, which suggests that high valuation is the dominant issue in judging the equity premium at this time.

It seems there is sufficient evidence in these markets, not only in their outward patterns but also in their correlation with each other and with other events, to feel pretty safe in concluding that we have seen a speculative bubble here. I know that there are

Figure 5. P/E for the S&P Composite in Relation to Subsequent 10-Year Real Composite Returns

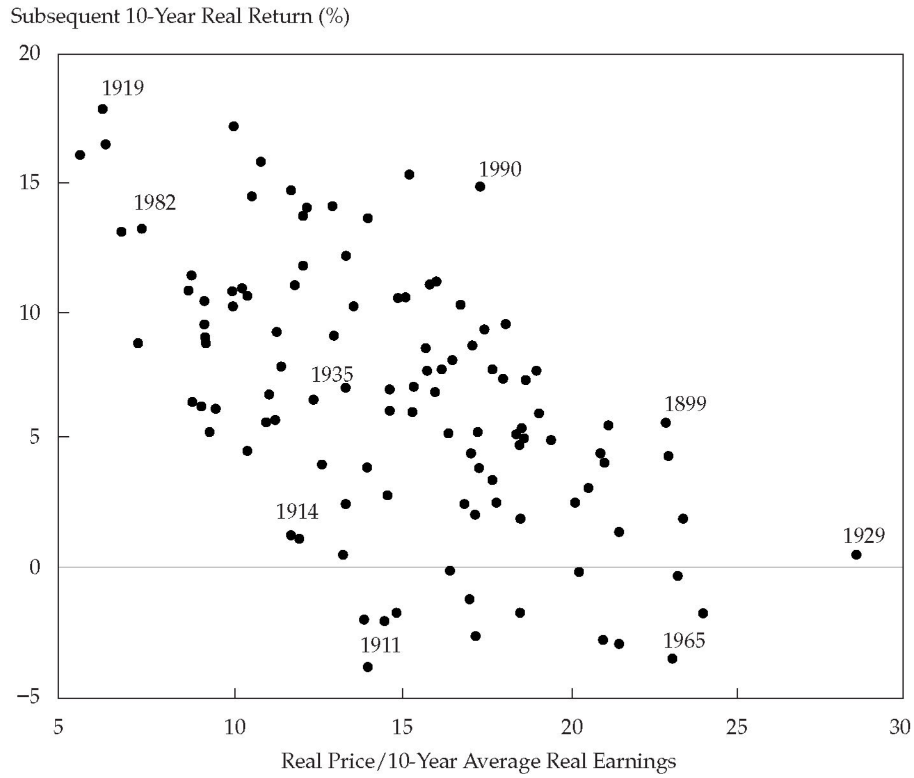

Notes: P/E for 1881–1990; average real returns for 1891–2000. A similar scattergram was used in the Campbell–Shiller presentation to Congress in 1996 (see Campbell and Shiller 1998) .

some academics who still apparently believe that there are no such things as speculative bubbles.[3] But these academics are increasingly in the minority in the profession.

Why Speculative Bubbles?

In *Irrational Exuberance*, I begin by showing the historical data that I just reviewed with you. The question that I addressed in the book is why we have speculative bubbles. I take three behavioral approaches to answering the question. In the first part, I consider structural factors—precipitating factors and amplification mechanisms—that encourage people to buy more stocks. The second part deals with cultural factors, such as the news media and "new era" theories. The third part deals with psychological factors, which include overconfidence, anchoring, and attention anomalies.

[3] For example, Peter Garber, in his recent (2000) book *Famous First Bubbles: The Fundamentals of Early Manias*, argues that even the tulipmania in Holland in the 1600s was essentially rational. He concludes, "The wonderful tales from the tulipmania are catnip irresistible to those with a taste for crying bubble, even when the stories are obviously untrue" (p. 83).

I have not heard many of these factors mentioned at our meeting today. It is puzzling to me that economists rarely seem to express an appreciation of the news media as important transmitters of speculative bubbles and of the idea that we are in a new era. Every time a speculative bubble occurs, many people who work in the media churn out stories that we are in a new era. I documented this phenomenon in my book by looking at a number of different cases in which the stock markets in various countries rose over a brief period, and I was able to find in each of them a new era theory in the newspaper.

Expert Theories

"Bubbles, Human Judgment, and Expert Opinion" was written to be of interest to practitioners. The objective was to observe how investors react to a market bubble and then try to interpret that phenomenon.

During the book tour for *Irrational Exuberance* in 2000 and 2001, I was often speaking to investment professionals, and although I had the sense that many times I was engaging their interest, I often did not have the sense that I was really connecting with them.

In many cases, they were not a really receptive audience. There was a sense of momentum or inertia among many of these people. They appeared to be of two minds—the one of an interested book reader and the other of a more rigid committee member or bureaucrat. I wanted to talk about that type of behavior in the "Bubbles" paper.

Why would that behavior be happening? What evidence would help us understand it? The reason I set forth in the paper is that the market is like a supertanker that cannot make sudden changes in course: Even if people like me present a case that the market is overpriced and is going to fall and even if people like me convince investment professionals that the market outlook is not so good, the professionals will not really make substantive changes in their portfolios. They may well continue to hold the 55 percent of their portfolios in U.S. equities and 11 percent in non-U.S. equities. University portfolio managers and other institutional investors were not withdrawing from the market in 1999.

In the paper, I discuss the *feedback* theory of bubbles that Andrei Shleifer and Nicholas Barberis (2000), I (1990), and others have talked about. In the feedback theory, demand for shares is modeled as a distributed lag of past returns plus the effect of precipitating factors. When returns have been high for a while, investors become more optimistic and bid up share prices, which amplifies the effects of precipitating factors. I consider this behavior to be an inconstancy in judgment, not naive extrapolation; for portfolio managers to respond naively to past returns seems implausible. Inconstancy in judgments arises because committees and their members find it difficult to respond accurately and incrementally to evidence, especially when the evidence is ambiguous, qualitative rather than quantitative, and ill defined. Ultimately, recent past returns have an impact on the decisions committee members make, even if they never change their conscious calculations. This feedback behavior thus amplifies the effect on the market of any precipitating factors that might initiate a speculative bubble.

The critical point is that the problem faced by institutional investors in deciding how much to put in the stock market is extremely complex; it has an infinite number of aspects that cannot possibly be completely analyzed. In such situations, people may fall into a pattern of behavior given by the "representative heuristic"—a psychological principle described by Kahneman and Tversky (1974, 1979) in which people tend to make decisions or judge information based on familiar patterns, preconceived categories or stereotypes of a situation. We tend to not take an objective outlook but to observe the similarity of a current pattern to a familiar, salient image in our minds and assume that the future will be like that familiar pattern.

Part of the problem that institutional investors face is the impossibility of processing all the available information. Ultimately, the decision whether to invest heavily in the stock market is a question of historical judgment. There are so many pieces of information that no one person can process all of them.

Therefore, institutional investment managers must rely on "conventional wisdom." They make decisions based on what they perceive is the generally accepted expert opinion. A problem with that approach is that one cannot know how much information others had in reaching the judgments laid out in conventional wisdom. In addition, investors do not know whether others were even relying on information or were, for their part, just using their judgment.

These kinds of errors that professionals make are analogous to the errors we sometimes make when, for example, we walk out of a conference and cross the street as a group. We may be talking about something interesting, so each person in the group assumes that someone else is looking at oncoming traffic. Sometimes, nobody is.

The tendency to follow conventional wisdom is increased by the strange standard we have called "the prudent person rule," part of fiduciary responsibility that is even written into ERISA. It is a strange standard because what it's really saying is not clear. As set forth in the ERISA regulations adopted in 1974, the prudent person rule states that investments must be made with

> the care, skill, and diligence, under the circumstances then prevailing, that a prudent man acting in a like capacity and familiar with such matters would use in the conduct of an enterprise with like character and like aims.

I interpret the statement to mean that an investment manager or plan sponsor must make judgments based on what is considered conventional at the time, not independent judgments.

The prudent person rule is a delicate attempt to legislate against stupidity, but the way the problem is addressed basically instructs the trustee or sponsor to be conventional. "Conventional" is exactly how I would describe what I think has happened to institutional investors and the way they approach the market. In 2000, many institutional investors believed they should not be so exposed to the market, but they could not justify to their organizations, within the confines of the prudent person rule, cutting back equity exposure. This dilemma is a serious problem.

Another problem that managers of institutional investments have can be described as "groupthink," a term coined in a wonderful book of the same name by the psychologist Irving Janis (1982). In the book, Janis gives case studies of committees or groups of highly intelligent people making big mistakes. In particular, he discusses the mistakes that arise because of group pressures individuals feel to conform. Janis points out that people who participate in erroneous decisions often find themselves censoring their statements because they believe, "If I express my dissenting view too often, I will be marginalized in the group and I will not be important." He uses the term "effectiveness trap" to describe this thinking. Dissenters, although they may be correct in their opinions, fear that they are likely to see their influence reduced if they express their opinions. Janis describes, for example, responses in the Lyndon Johnson administration to a Vietnam bombing fiasco. When Johnson wrote about this episode in his memoirs, he did not mention any substantial dissent. Yet, those involved remember having dissenting views. Evidently, they did not express their views in such a way that Johnson remembered the dissent after the fact.

As economists, we talk a great deal about models, which concretize the factors in decisions, but when you are making a judgment about how to manage a portfolio, you face real-world situations. The real world is fundamentally uncertain. And fundamental uncertainty is what Knight talks about in *Risk, Uncertainty and Profit* (1964): How do we react in committees or as groups or as individuals within groups?

An argument Shafir, Simonson, and Tversky (2000) recently made that they applied to individual decisions is, I think, even more applicable to group decisions. The authors stated that when we are making what seems like a portentous decision, our minds seek a *personalized* way to justify the decision; we do not simply consider what to do. They asked people to make hypothetical custody decisions about divorcing couples. They described the two parents and then asked each participant to choose which parent would get custody of the child. They framed the question in two different ways. One question was, "Which parent would you give the child to?" And the other was, "Which parent would you deny custody to?" Of course, the question is the same either way it is framed. Nevertheless, the authors found systematic differences in the responses. When the parents were described, one person was described in bland terms and the other person in very vivid terms—both good extremes and bad extremes. Participants tended to point their decisions to the more salient person (the more vividly described person) in the couple. For example, when the question was framed for awarding custody, participants tended to award custody to the person who was vividly described—even though the description included bad things. And when the question was framed for denying custody, participants tended to deny custody to the person who was vividly described—even though the description included good things.

This research points to a fundamental reason for inertia in organizations: Institutions have to have a very good reason to change any long-standing policy, but the kinds of arguments that would provide that good reason are too complicated (not salient enough) to be persuasive.

Conclusion

My talk has taken us a little bit away from the abstract issue of the long-run equity premium that has been talked about so much at this forum. I have described a shorter-run phenomenon, the recent stock market bubble, and I have described some particular psychological principles that must be borne in mind if we are to understand this recent behavior. But we cannot see the weaknesses of faulty abstract principles unless we focus on particular applications of the principles. I hope that my discussion today has raised issues relevant to understanding whether we ought to consider the markets efficient, whether we ought to be "puzzled" by the past equity premium, and whether we should expect this historical premium to continue in the future.

Current Estimates and Prospects for Change I

Robert J. Shiller
Yale University
New Haven, Connecticut

SUMMARY

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

Robert Shiller described the equity premium puzzle as inseparable from the foundations of behavioral finance. The three bases of his presentation were

- Campbell and Shiller, testimony before the Federal Reserve Board on December 3, 1996,[1]
- *Irrational Exuberance* (published in April 2000; see Shiller 2000), and
- "Bubbles, Human Judgment, and Expert Opinion" (Shiller 2002).

[1] Summarized in Campbell and Shiller (1998).

The third publication was aimed at (nonprofit) practitioners (particularly, those at U.S. educational endowments). Much behavioral finance describes apparently foolish behavior in the market, but trustees are, presumably, intelligent people. Yet, even they have not been betting against the market during the recent bubble. Despite warnings, intelligent people have not lost faith in the stock market. Why is expert opinion so biased?

Shiller's **Figure 1** showed the real Nasdaq Composite Index from October 1984 to October 2001. It provided clear evidence of a perfect bubble from 1999 to late 2000. The same could be seen in his **Figure 2** of the S&P Composite Index from 1871 to 2001. Two other, lesser bubbles appeared—in the late 1920s and the late 1960s. Similarly, the **Figure 3** graph of real estate prices in Los Angeles, California, showed a clear bubble (although it was smoother than the market bubble) around 1990. **Figure 4**, of the S&P

Figure 1. Real Nasdaq Composite, October 1984–October 2001

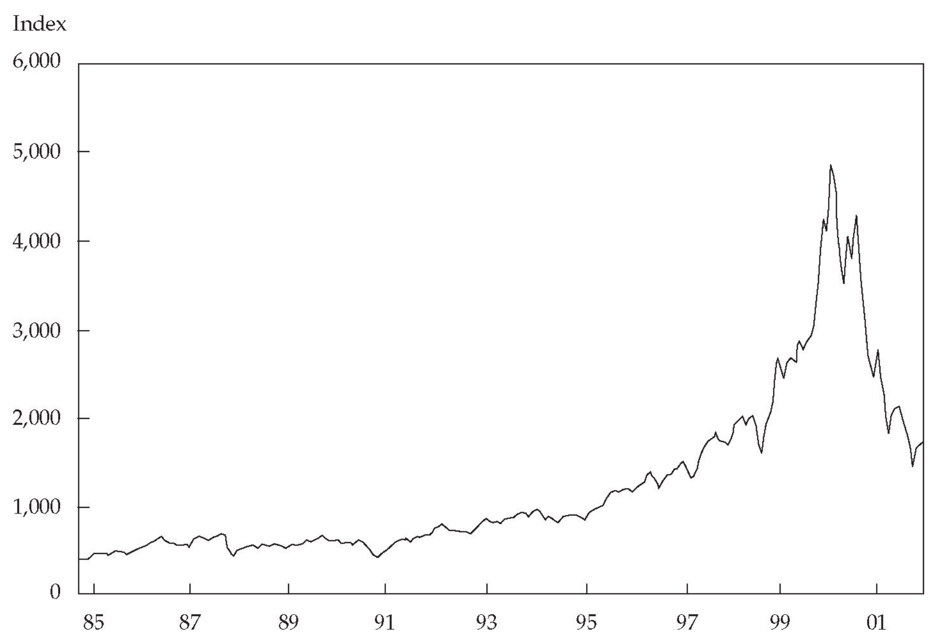

Figure 2. S&P Composite: Real Price and Earnings, January 1871–2001

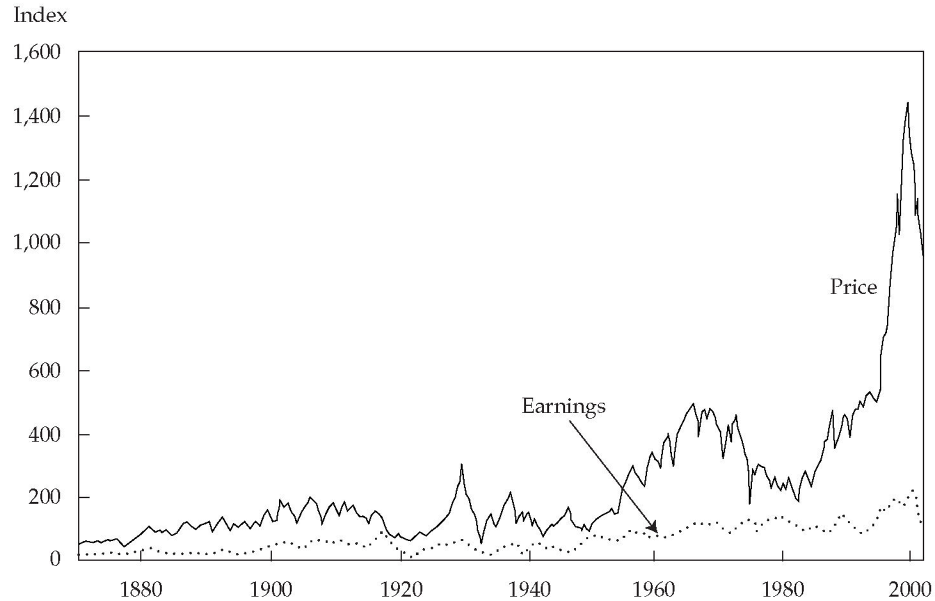

Note: Measured monthly.

Figure 3. Case–Shiller Home Price Index: Los Angeles Single-Family Home Prices, Fourth Quarter 1975 to Second Quarter 2001

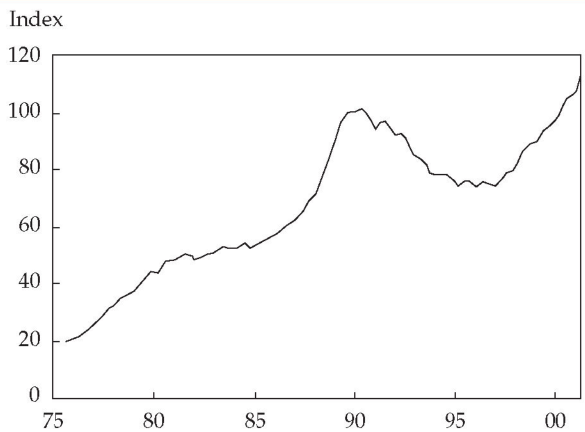

Note: Measured quarterly.

Composite P/E (real price divided by average real earnings over the preceding 10 years) from 1881 to 2001, showed bubbles recently, in the late 1920s, around 1900 (to a lesser extent), in the late 1930s, and in the 1960s.

Figure 5 is a scattergram showing how the S&P Composite P/E predicts future S&P Composite returns. The P/E is now around the 1929 level, which suggests that valuation is the dominant issue in terms of the equity premium at this time.

Figure 4. P/E for the S&P Composite, January 1881–October 2001

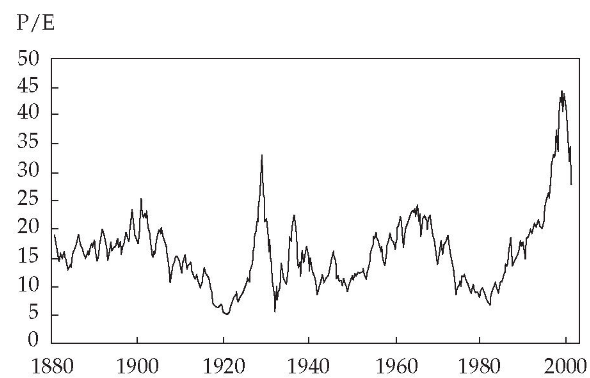

Note: P/E calculated as price over 10-year lagging earnings (a calculation recommended by Graham and Dodd in 1934).

In his book *Irrational Exuberance*, Shiller dealt with three types of factors leading to excessive valuations: structural, cultural, and psychological. Cultural factors included the news media and "new era" theories. The news media are important transmitters of speculative bubbles, and every bubble is accompanied by a new era theory to explain the rise in prices. Among psychological factors are overconfidence, anchoring, and attention anomalies.

Figure 5. P/E for the S&P Composite in Relation to Subsequent 10-Year Real Composite Returns

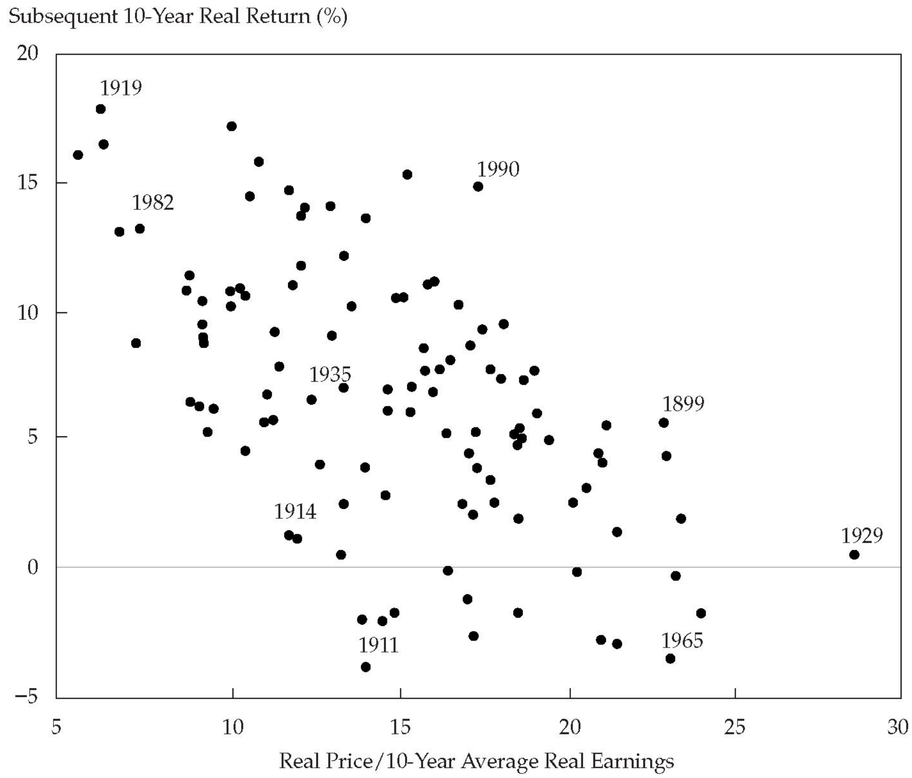

Notes: P/E for 1881–1990; average real returns for 1891–2000. A similar scattergram was used in the Campbell–Shiller presentation to Congress in 1996 (see Campbell and Shiller 1998) .

Turning to the subject of his "Bubbles" paper, Shiller discussed a number of aspects of behavioral finance behind the behavior of investment professionals that drove equity prices up. The most important factor is the inertia of a bureaucratic process. No matter how convincing the evidence that stock prices are too high, institutional committees do not change their asset allocations, which were generally about 60 percent in U.S. and non-U.S. equities in 1999.

The influence of recent past returns is powerful. Reliance on recent returns might be thought of as naive extrapolation, but Shiller prefers to think of it as inconstancy in judgment. It is difficult for committees to maintain the same judgment at all times when the evidence is ambiguous and complicated. The tendency is to assume that the future will be like the past.

The impossibility of processing all available information leads to reliance on conventional wisdom. Institutional investors have a tendency to trust the opinions of others without knowing what information those others are making use of. Moreover, the "prudent person rule" is, unfortunately, to "do what is conventional."

Shiller also cited examples of the "effectiveness trap"—the group pressure to conform—described in *Groupthink* (Janis 1982). Dissenters, although they may be correct in their opinions, fear that they are likely to see their influence reduced if they express their opinions. Other references Shiller made dealt with the difficulty of getting organizations to change long-standing policy. Committees need a *very* good reason to change a policy.

Shiller's conclusions included the following:

- Bubble behavior and the equity risk premium are tied up with many issues of human cognition and judgment.
- Institutional investors have generally been too slow to react to the negative equity premium today.

Current Estimates and Prospects for Change II

Rajnish Mehra

Professor of Finance
University of California, Santa Barbara
National Bureau of Economic Research and Vega Asset Management

Analysts have more than 100 years of good, clean economic data on asset returns that support the persistence of a historical long-term U.S. equity risk premium over U.S. T-bills of 5–7 percent (500–700 bps)—but the expected equity risk premium an analyst might have forecasted at the beginning of this long period was about 2 percent. The puzzle is that stocks are not so much riskier than T-bills that a 5–7 percent difference in rates of return is justified. Analyses of the long series of data indicate that the relationship between *ex ante* and *ex post* premiums is inverse. The relationship between the market and the risk premium is also inverse: When the value of the market has been high, the mean equity risk premium has been low, and vice versa. Finally, investors and advisors need to realize that all conclusions about the equity risk premium are based on and apply only to the very long term. To predict next year's premium is as impossible as predicting next year's stock returns.

I took the topic of the equity risk premium literally and considered, given current valuation levels, what is the expected equity risk premium. I would argue that this question is an exercise in forecasting and has little to do with the academic debate on whether the historically observed equity risk premium has been a puzzle. Let me illustrate.

Table 1 shows the data available to us from various sources and research papers on U.S. equity returns (generally proxied by a broad-based stock index), returns to a relatively riskless security (typically a U.S. Treasury instrument), and the equity risk premium for various time periods since 1802. The equity premium can be different over the same time period, primarily because some researchers measure the premium relative to U.S. T-bonds and some measure it relative to T-bills. The original Mehra–Prescott paper (1985) measured the premium relative to T-bills. Capital comes in a continuum of risk types, but aggregate capital stock in the United States will give you a return of about 4 percent. If you combine the least risky part and the riskier part, such as stocks, their returns will be different but will average about 4 percent. I can, at any time, pry off a very risky slice of the capital risk continuum and compare its rate of return with another slice of the capital risk continuum that is not at all risky.

Table 1 provides results from a fairly long series of data—almost 200 years—and the premium exists even when the bull market between 1982 and 2000 is

Table 1. Real U.S. Equity Market and Riskless Security Returns and Equity Risk Premium, 1802–2000

Period	Mean Real Return on Market Index	Mean Real Return on Relatively Riskless Asset	Risk Premium
1802–1998	7.0%	2.9%	4.1%
1889–2000	7.9	1.0	6.9
1889–1978	7.0[a]	0.8	6.2[b]
1926–2000	8.7	0.7	8.0
1947–2000	8.4	0.6	7.8

[a]Not rounded, 6.98 percent.
[b]Not rounded, 6.18 percent.

Sources: Data for 1802–1998 are from Siegel (1998); for 1889–2000, from Mehra and Prescott (1985).

excluded. That bull market certainly contributed to the premium, but the premium is pretty much the same in all the periods. One comment on early-19th-century data: The reason Edward Prescott and I began at 1889 in our original study is that the earlier data are fairly unreliable. The distinction between debt and equity prior to 1889 is fuzzy. What was in a basket of stocks at that time? Would bonds actually be called risk free? Because the distinction between these types of capital was unclear, the equity premium for the 1802–1998 period appears to be lower in Table 1 than I believe it really was. As **Table 2** shows, the existence of an equity premium is consistent across developed countries—at least for the post-World War II period.

The puzzle is that, adjusted for inflation, the average annual return in the U.S. stock market over 110 years (1889–2000) has been a healthy 7.9 percent, compared with the 1 percent return on a relatively riskless security. Thus, the equity premium over that time period was a substantial 6.2 percent (620 basis points). One could dismiss this result as a statistical artifact, but those data are as good an economic time series as we have. And if we assume some stationarity in the world, we should take seriously numbers that show consistency for 110 years. If such results occurred only for a couple of years, that would be a different story.

Is the Premium for Bearing Risk?

This puzzle defies easy explanation in standard asset-pricing models. Why have stocks been such an attractive investment relative to bonds? Why has the rate of return on stocks been higher than on relatively risk-free assets? One intuitive answer is that because stocks are "riskier" than bonds, investors require a larger premium for bearing this additional risk; and indeed, the standard deviation of the returns to stocks (about 20 percent a year historically) is larger than that of the returns to T-bills (about 4 percent a year). So, obviously, stocks are considerably more risky than bills!

But are they?

Why do different assets yield different rates of return? Why would you expect stocks to give you a higher return? The deus ex machina of this theory is that assets are priced such that, *ex ante*, the loss in marginal utility incurred by sacrificing current consumption and buying an asset at a certain price is equal to the expected gain in marginal utility contingent on the anticipated increase in consumption when the asset pays off in the future.

The operative emphasis here is the *incremental loss or gain* of well-being resulting from consumption, which should be differentiated from incremental consumption because the same amount of consumption may result in different degrees of well-being at different times. (A five-course dinner after a heavy lunch yields considerably less satisfaction than a similar dinner when one is hungry!)

As a consequence, assets that pay off when times are good and consumption levels are high—that is, when the incremental value of additional consumption is low—are less desirable than those that pay off an equivalent amount when times are bad and additional consumption is both desirable and more highly valued.

Let me illustrate this principle in the context of a popular standard paradigm, the capital asset pricing model (CAPM). This model postulates a linear relationship between an asset's "beta" (a measure of systematic risk) and expected return. Thus, high-beta stocks yield a high expected rate of return. The reason is that in the CAPM, good times and bad times are captured by the return on the market. The performance of the market as captured by a broad-based index acts as a surrogate indicator for the relevant state of the economy. A high-beta security tends to pay off more when the market return is high, that is, when times are good and consumption is plentiful; as

Table 2. Real Equity and Riskless Security Returns and Equity Risk Premium: Selected Developed Markets, 1947–98

Country	Period	Mean Real Return on Market Index	Mean Real Return on Relatively Riskless Asset	Risk Premium
United Kingdom	1947–1999	5.7%	1.1%	4.6%
Japan	1970–1999	4.7	1.4	3.3
Germany	1978–1997	9.8	3.2	6.6
France	1973–1998	9.0	2.7	6.3

Sources: Data for the United Kingdom are from Siegel (1998); the remaining data are from Campbell (2002).

discussed earlier, such a security provides less incremental utility than a security that pays off when consumption is low, is less valuable to investors, and consequently, sells for less. Thus, assets that pay off in states of low marginal utility will sell for a lower price than similar assets that pay off in states of high marginal utility. Because rates of return are inversely proportional to asset prices, the latter class of assets will, on average, give a lower rate of return than the former.

Another perspective on asset pricing emphasizes that economic agents prefer to smooth patterns of consumption over time. Assets that pay off a relatively larger amount at times when consumption is already high "destabilize" these patterns of consumption, whereas assets that pay off when consumption levels are low "smooth" out consumption. Naturally, the latter are more valuable and thus require a lower rate of return to induce investors to hold them. (Insurance policies are a classic example of assets that smooth consumption. Individuals willingly purchase and hold them in spite of their very low rates of return.)

To return to the original question: Are stocks that much riskier than bills so as to justify a 7 percent differential in their rates of return?

What came as a surprise to many economists and researchers in finance was the conclusion of a research paper that Prescott and I wrote in 1979. Stocks and bonds pay off in approximately the same states of nature or economic scenarios; hence, as argued earlier, they should command approximately the same rate of return. In fact, using standard theory to estimate risk-adjusted returns, we found that stocks on average should command, at most, a 1 percent return premium over bills. Because for as long as we had reliable data (about 100 years), the mean premium on stocks over bills was considerably and consistently higher, we realized that we had a puzzle on our hands. It took us six more years to convince a skeptical profession and for our paper (the Mehra and Prescott 1985 paper) to be published.

Ex Post versus *Ex Ante*

Some academicians and professionals hold the view that at present, there is no equity premium and, by implication, no equity premium puzzle. To address these claims, we need to differentiate between two interpretations of the term "equity premium." One interpretation is the *ex post* or realized equity premium over long periods of time. It is the actual, historically observed difference between the return on the market, as captured by a stock index, and the risk-free rate, as proxied by the return on T-bills.

The other definition of the equity premium is the *ex ante* equity premium—a forward-looking measure. It is the equity premium that is *expected* to prevail in the future or the conditional equity premium given the current state of the economy. I would argue that it *must* be positive because all stocks must be held.

The relationship between *ex ante* and *ex post* premiums is inverse. After a bull market, when stock valuations are exceedingly high, the *ex ante* premium is likely to be low, and this is precisely the time when the *ex post* premium is likely to be high. After a major downward correction, the *ex ante* (expected) premium is likely to be high and the realized premium will be low. This relationship should not come as a surprise because returns to stock have been documented to be mean reverting. Over the long term, the high and low premiums will average out.

Which of these interpretations of the equity risk premium is relevant for an investment advisor? Clearly, the answer depends on the planning horizon.

The historical equity premium that Prescott and I addressed in 1985 is the premium for very long investment horizons, 50–100 years. And it has little—in fact, nothing—to do with what the premium is going to be over the next couple of years. Nobody can tell you that you are going to get a 7 percent or 3 percent or 0 percent premium next year.

The *ex post* equity premium is the realization of a stochastic process over a certain period, and as **Figure 1** shows, it has varied considerably over time. Furthermore, the variation depends on the time horizon over which it is measured. Over this 1926–2000 period, the realized equity risk premium has been positive and it has been negative; in fact, it has bounced all over the place. What else would you expect from a stochastic process in which the mean is 6 percent and the standard deviation is 20 percent? Now, note the pattern for 20-year holding periods in **Figure 2**. This pattern is more in tune with what Jeremy Siegel was talking about [see the "Historical Results" session]. You can see that over 20-year holding periods, there is a nice, decent premium.

Figure 3 carries out exactly the exercise that Brad Cornell recommended [see the "Historical Results" session]: It looks at stock market value (MV)—that is, the value of all the equity in the United States—as a share of National Income (NI). These series are cointegrated, so when you divide one by the other, you get a stationary process. The ratio has been as high as approximately 2 times NI and as low as approximately 0.5 NI. The graph in Figure 3 represents risk. If you are looking for stock market risk, you are staring at it right here in Figure 3. This risk is low-frequency, persistent risk, not the year-to-year volatility in the market. This persistence defies easy

Figure 1. Realized Equity Risk Premium per Year, January 1926–January 2000

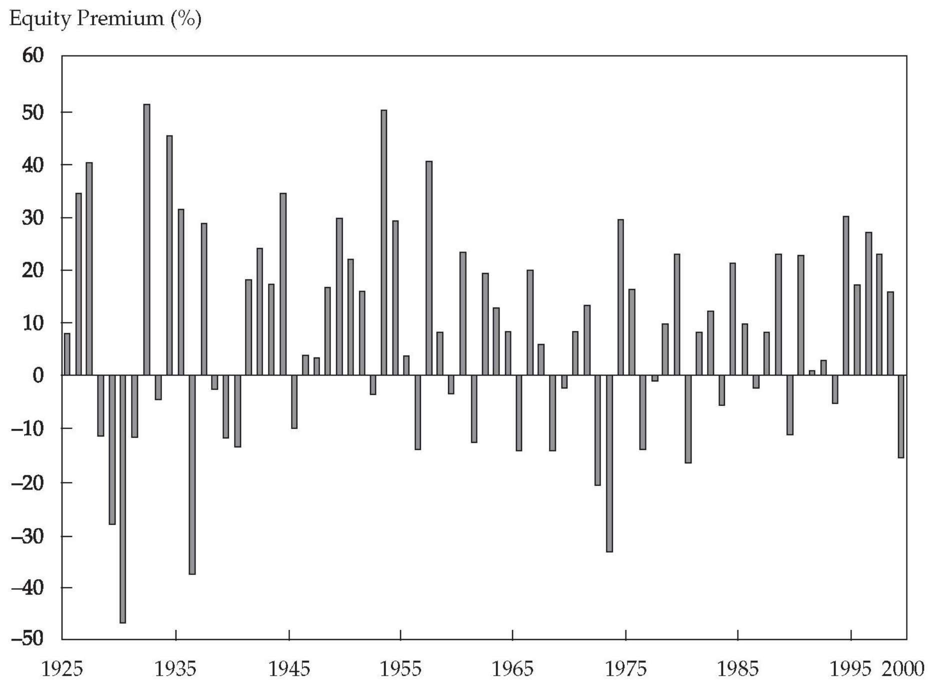

Source: Ibbotson Associates (2001).

Figure 2. Mean Equity Risk Premium by 20-Year Holding Periods, January 1926–January 2000

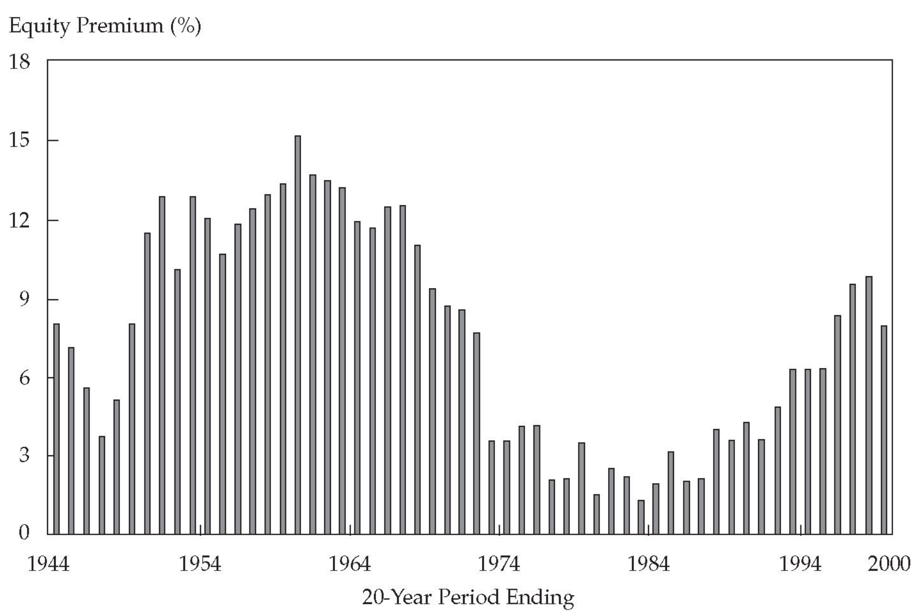

Source: Ibbotson Associates (2001).

Figure 3. U.S. Stock Market Value/National Income, January 1929–January 2000

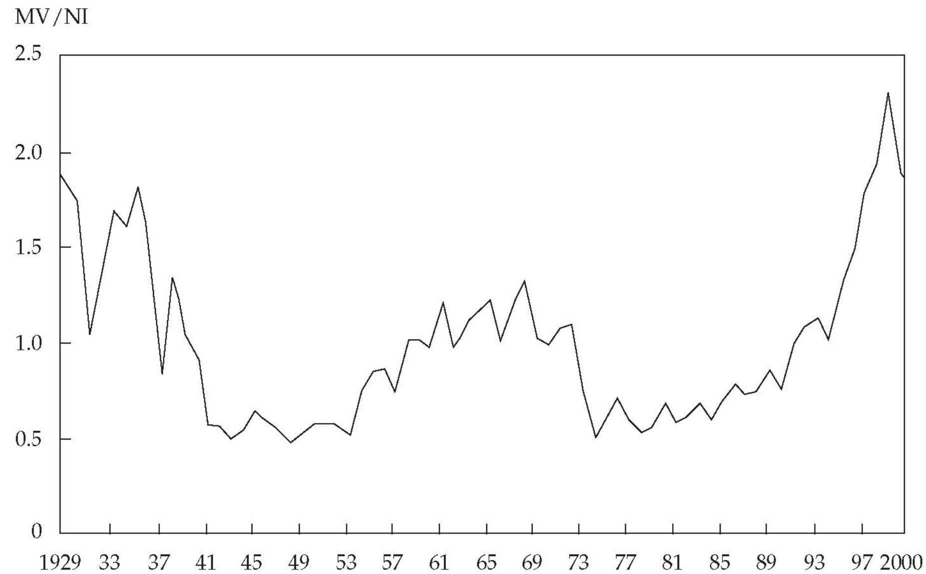

Source: Data updated from Mehra (1998).

explanation for the simple reason that if you look at cash flows over the same period of time relative to GDP, they are almost trendless. There are periods of relative overvaluation and periods of undervaluation, and they seem to persist over time.

When I plotted the contemporaneous equity risk premium over the same period, the graph I got was not very informative, so I arbitrarily broke up the data into periods when the market was more than 1 NI and when the market was below 1 NI. I averaged out all the wiggles in the equity premium graph, and **Figure 4** shows the smoothed line overlaid on the graph from Figure 3 of MV/NI. As you can see, when the market was high, the mean equity risk premium was low, and when the market was low, the premium was high.

Figure 4. Mean Equity Risk Premium and Market Value/National Income, January 1929–January 2000

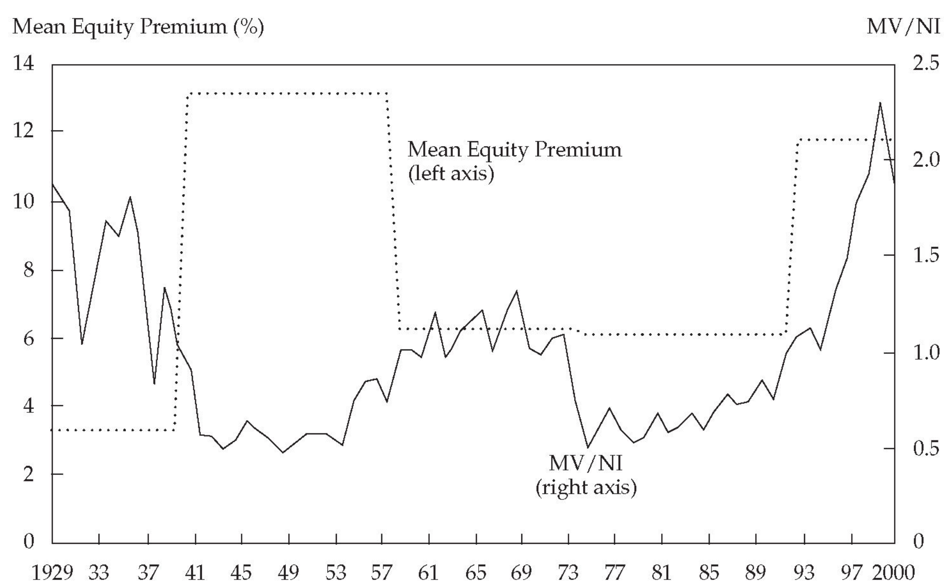

The mean equity risk premium three years ahead is overlaid on the graph of market value to net income in **Figure 5**. (The premium corresponding to 1929 on the dotted line represents the mean equity risk premium averaged from 1929 to 1932. So, the premium line ends three years before 2001). You can clearly see that the mean equity risk premium is much higher when valuation levels are low.

I might add that the MV/NI graph is the basis of most of the work in finance on predicting returns based on price-to-dividends ratios and price-to-earnings ratios. Essentially, we have historical data for only about two cycles. Yet, a huge amount of research and literature is based on regressions run with only these data.

A scatter diagram of MV/NI versus the mean three-year-ahead equity risk premium is shown in **Figure 6**. Not much predictability exists, but the relationship is negative. (The graphs and scatter diagrams for a similar approach but with the equity risk premium five years ahead are similar).

Finally, **Figure 7** plots mean MV/NI versus the mean equity risk premium three years ahead, but I arbitrarily divided the time into periods when MV/NI was greater than 1 and periods when it was less than 1, and I averaged the premium over the periods. This approach shows, on average, some predictability: Returns are higher when markets are low relative to GDP. But if I try to predict the equity premium over a year, for example, the noise dominates the drift.

Operationally, because the volatility of market returns is 20 percent, you do not get much information from knowing that the mean equity premium is 2 percent rather than 6 percent. From an asset-allocation point of view, I doubt that such knowledge would make any difference over a short time horizon—the next one or two years. The only approach that makes sense in this type of analysis is to estimate the equity premium over the very long horizon. The problem of predicting the premium in the short run is as difficult as predicting equity returns in the short run. Even if the conditional equity premium given current market conditions is small (and the general consensus is that it is), that fact, in itself, does not imply either that the historical premium was too high or that the unconditional equity premium has diminished.

Looking into the Future

If this analysis had been done in 1928, what would an exercise similar to what Prescott and I did in 1985 have yielded? Suppose the analysis were done for the period from 1889 to 1928; in 1929, the mean real return on the S&P 500 was 8.52 percent, the mean real return on risk-free assets was 2.77 percent, and thus the observed mean equity premium would have been 5.75 percent. A theoretical analysis similar to Prescott's and mine would have yielded a 2 percent equity premium.

Figure 5. Mean Equity Risk Premium Three Years Ahead and Market Value/National Income, January 1929–January 2000

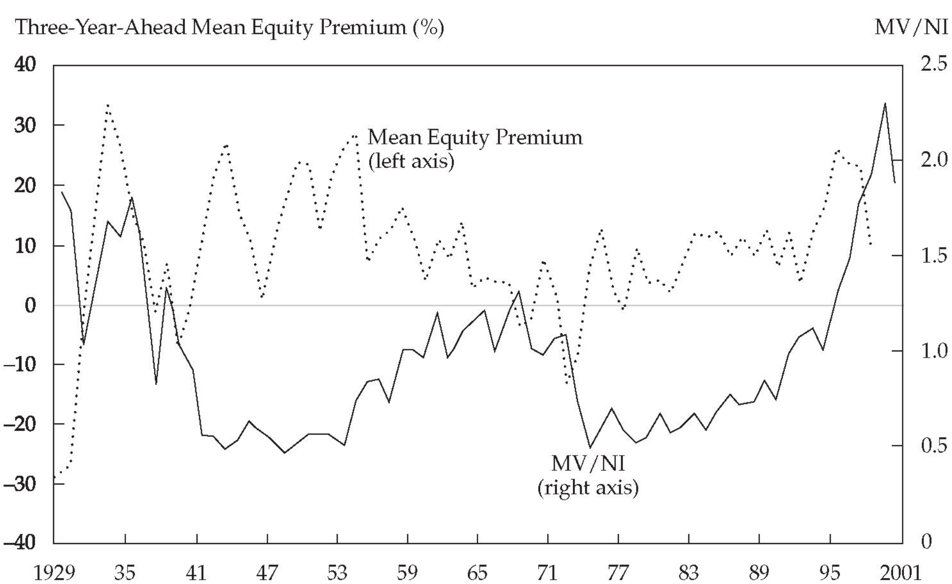

Figure 6. Scatter Diagram: Mean Equity Risk Premium Three Years Ahead versus Market Value/National Income, January 1929–January 2000 Data

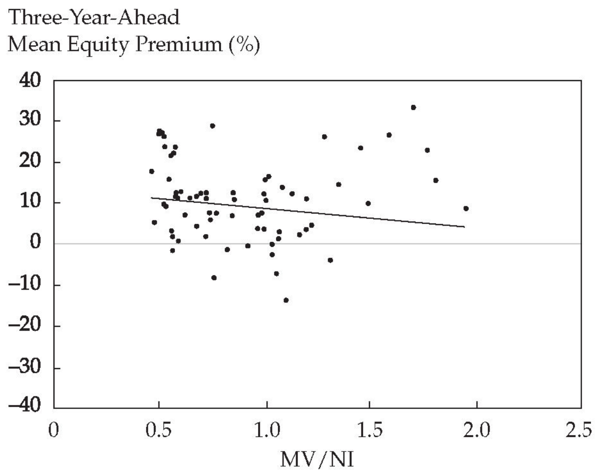

Note: $y = 4.7159x + 13.321$.

What could have been concluded from that information? The premium of 2 percent is the realization of a stochastic process with a large standard deviation. If the investor of 1928 saw any pattern in the stochastic process, optimizing agents would have endogenously changed the prices. That understanding makes it much more difficult to say we have a bubble. What we see is only one realization of a stochastic process. We would ideally like to see the realizations in many different, parallel universes and see how many times we actually came up with 2 percent and how many times we didn't. However, we are constrained by reality and observe only one realization!

The data used to document the equity premium are as good and clean as any economic data that I have seen. A hundred years of economic data is a long time series. Before we dismiss the equity premium, not only do we need to understand the observed phenomena (why an equity risk premium should exist), but we also need a plausible explanation as to why the future is likely to be different from the past. What factors may be important in determining the future premium? Life-cycle and demographic issues may be important, for example; the retirement of aging Baby Boomers may cause asset deflation. If so, then the *realized* equity premium will be low in 2010. But if asset valuations are expected to be low in 2010, why should the premium not be lower now? Perhaps what we are seeing in the current economy is the result of market efficiency taking the aging Baby Boomers into account. Either we will understand why a premium should exist (in which case, it will persist), or if it is a statistical artifact, it should disappear now that economic agents are aware of the phenomenon.

Figure 7. Mean Equity Risk Premium Three Years Ahead by Time Periods and Market Value/National Income, January 1929–January 2000

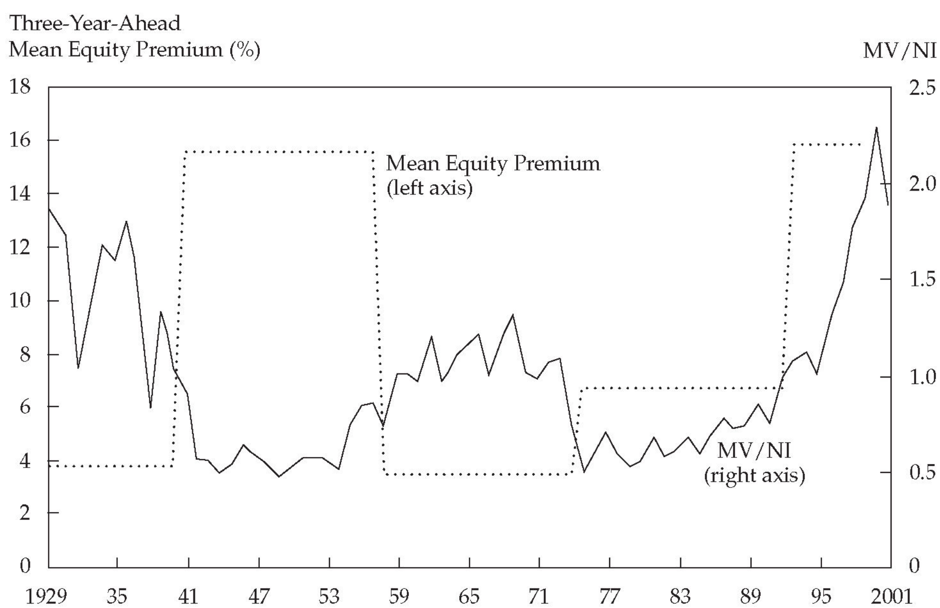

Note: The equity premium was averaged over time periods in which MV/NI > 1 and MV/NI < 1.

Current Estimates and Prospects for Change II

Rajnish Mehra
Professor of Finance
University of California, Santa Barbara
National Bureau of Economic Research and Vega Asset Management

SUMMARY

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

Rajnish Mehra proposed that analyzing the equity risk premium is an exercise in forecasting that has little to with the academic debate over whether the observed past excess return on equities presents a puzzle. Why is the equity premium a puzzle?

Table 1 shows real returns for long and not-so-long periods of time for the U.S. stock market, a relatively riskless asset, and the risk premium. A real return on equities of about 7 percent characterizes some long time periods, including 1889–1978, a period that did not incorporate the recent bull market. For the 1889–2000 period, the return was 7.9 percent. The standard deviation of annual returns was about 20 percent. Moreover, as **Table 2** shows, other countries have shown similar returns.

Table 1. Real U.S. Equity Market and Riskless Security Returns and Equity Risk Premium, 1802–2000

Period	Mean Real Return on Market Index	Mean Real Return on Relatively Riskless Asset	Risk Premium
1802–1998	7.0%	2.9%	4.1%
1889–2000	7.9	1.0	6.9
1889–1978	7.0[a]	0.8	6.2[b]
1926–2000	8.7	0.7	8.0
1947–2000	8.4	0.6	7.8

[a]Not rounded, 6.98 percent.
[b]Not rounded, 6.18 percent.

Sources: Data for 1802–1998 are from Siegel (1998); for 1889–2000, from Mehra and Prescott (1985).

U.S. T-bills have returned about 1 percent with a 4 percent standard deviation. Why are the returns on T-bills so different from those on equity? We might say we are looking at an aberration, but this time series is the best evidence we have. The difference defies easy explanation by standard asset-pricing models. Is it explained by risk differences? The answer is not clear.

Our theory tells us that assets are priced in such a way that, *ex ante*, the loss in marginal utility incurred by sacrificing current consumption to buy an asset at a certain price is equal to the expected gain in marginal utility contingent on the anticipated increase in consumption when the asset pays off in the future. The emphasis here is on *incremental loss or gain* of utility of consumption, which should be differentiated from incremental consumption because the same amount of consumption may result

Table 2. Real Equity and Riskless Security Returns and Equity Risk Premium: Selected Developed Markets, 1947–98

Country	Period	Mean Real Return on Market Index	Mean Real Return on Relatively Riskless Asset	Risk Premium
United Kingdom	1947–1999	5.7%	1.1%	4.6%
Japan	1970–1999	4.7	1.4	3.3
Germany	1978–1997	9.8	3.2	6.6
France	1973–1998	9.0	2.7	6.3

Sources: Data for the United Kingdom are from Siegel (1998); the remaining data are from Campbell (2002).

in different degrees of well-being at different times. As a consequence, assets that pay off when times are good and consumption levels are high—i.e., when the marginal utility of consumption is low—are less desirable than those that pay off an equivalent amount when times are bad and additional consumption is more highly valued.

This theory is readily illustrated in the context of the capital asset pricing model, in which good times and bad times are captured by the return on the market. Why do high-beta stocks yield a high expected rate of return? A high-beta security tends to pay off more when the market return is high—that is, when times are good and consumption is plentiful. Such a security provides less incremental utility than a security that pays off when consumption is low, is less valuable, and consequently, sells for less. Because rates of return are inversely proportional to asset prices, the former class of assets will, on average, give a higher rate of return than the latter.

Another perspective emphasizes that economic agents prefer to smooth patterns of consumption over time. Assets that pay off a relatively larger amount at times when consumption is already high "destabilize" these patterns of consumption, whereas assets that pay off when consumption levels are low "smooth" out consumption. Naturally, the latter are more valuable and thus require a lower rate of return to induce investors to hold them. And such assets are purchased despite their very low expected rates of return. Insurance is an example.

What is surprising is that stocks and bonds *pay off in approximately the same states of nature* or economic scenarios. Hence, as Mehra argued earlier, they should command approximately the same rate of return. Using standard theory to estimate risk-adjusted returns, Mehra and Prescott (1985) showed that stocks, on average, should command, at most, a 1 percent (100 bps) return premium over bills. This finding presented a puzzle because the historically observed mean premium on stocks over bills was considerably and consistently higher.

The *ex post* excess return has varied a lot, which is not surprising. Graphs of the annual realized excess return in **Figure 1** and of the excess return for 20-year periods in **Figure 2** show dramatic differences.

Mehra stressed that we need to distinguish the *ex post* excess return on equity from the *ex ante* risk premium. The expected equity premium *must* be positive. Following a bull market, the *ex post* will be high and the *ex ante* will be low. Over time, they will average out. A conclusion for the future depends on the planning horizon. Mehra was addressing the premium for the very long term—on the order of 50–100 years. In the short term, as in Figure 1, the variance in returns makes it quite impossible to come up with any reliable forecast. Figure 2 for 20-year periods, however, shows something more promising.

Figure 1. Realized Equity Risk Premium per Year, January 1926–January 2000

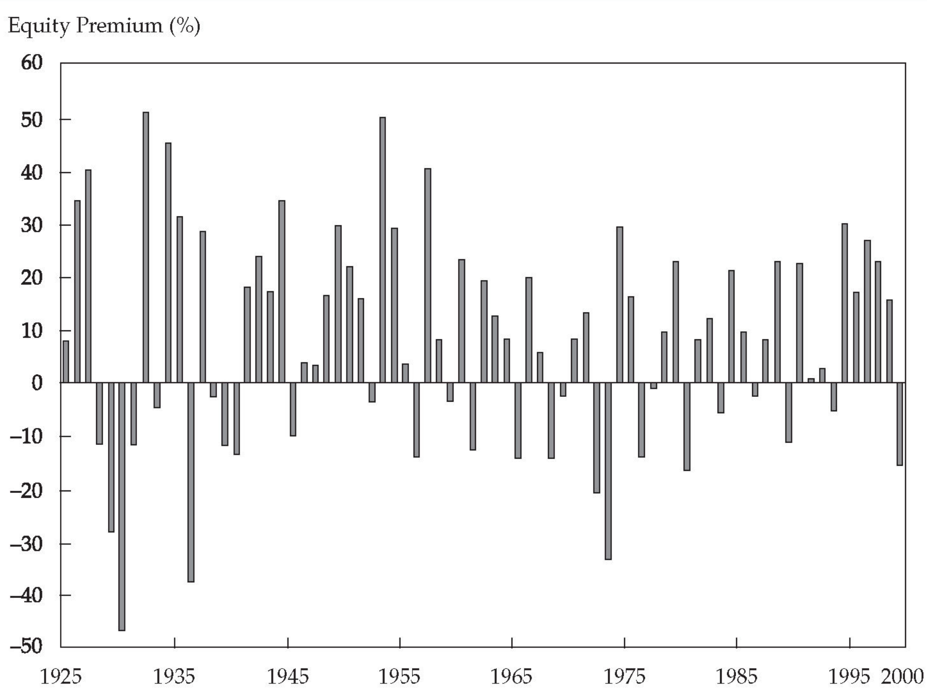

Source: Ibbotson Associates (2001).

Figure 2. Mean Equity Risk Premium by 20-Year Holding Periods, January 1926–January 2000

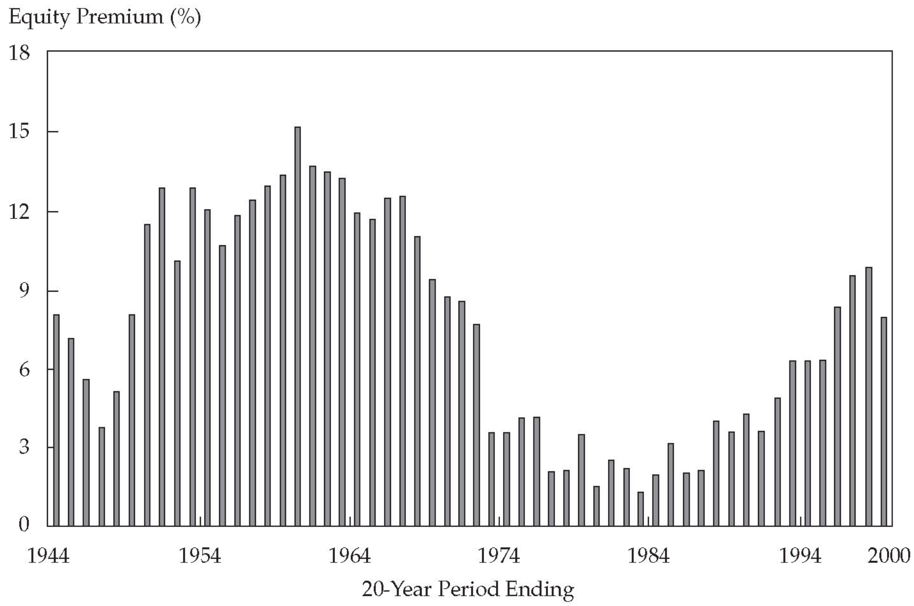

Source: Ibbotson Associates (2001).

Mehra's **Figure 3** showed the ratio of market value of equity (MV) to national income (NI) since 1929, and his **Figure 5** overlaid on that graph the three-year-ahead equity premium.[1] The ratio has ranged from 2 × NI to 0.5 × NI to 2.25 × NI. In **Figure 7**, Mehra split the 1929–2000 period into subperiods—those in which MV as a ratio of NI was greater than 1 and those in which it was less than 1—and overlaid on that graph is the three-year-ahead mean equity premium. Figure 7 shows that we have had two and a half cycles since 1929, and they reveal some predictive ability: On average, when MV/NI is low, the risk premium is high, which is useful as a guide for the very long term.

[1] Table and figure numbers in each Summary correspond to the table and figure numbers in the full presentation.

Figure 3. U.S. Stock Market Value/National Income, January 1929–January 2000

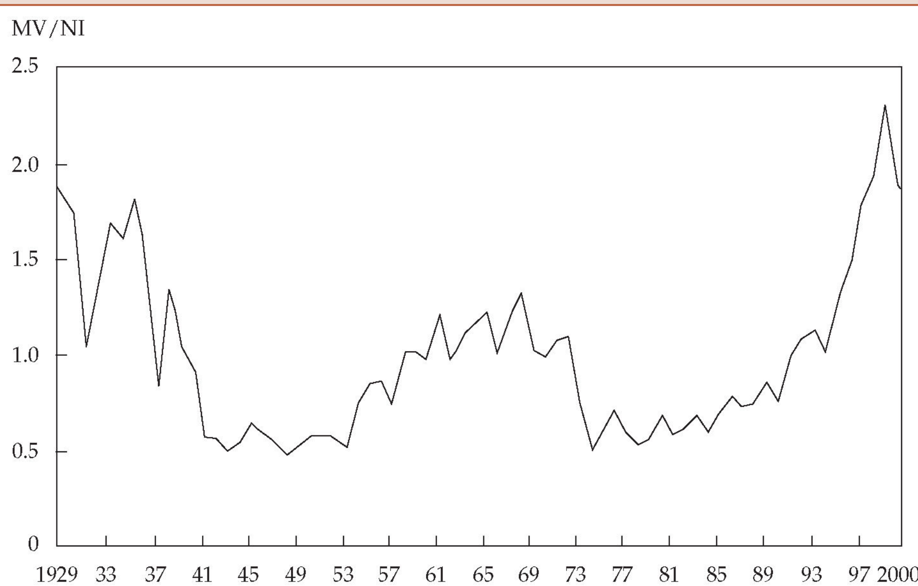

Source: Data updated from Mehra (1998).

Figure 5. Mean Equity Risk Premium Three Years Ahead and Market Value/National Income, January 1929–January 2000

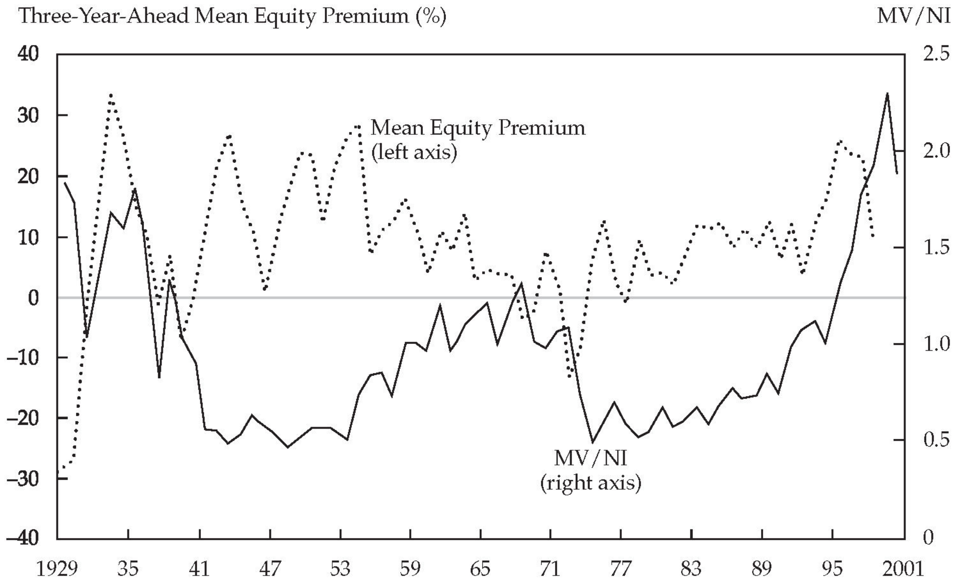

Figure 7. Mean Equity Risk Premium Three Years Ahead by Time Periods and Market Value/National Income, January 1929–January 2000

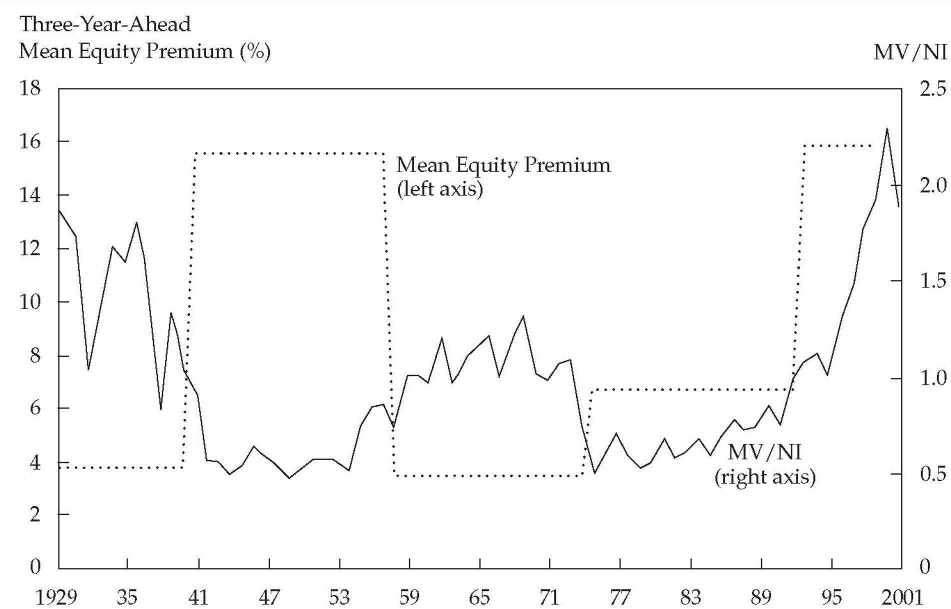

Note: The equity premium was averaged over time periods in which MV/NI > 1 and MV/NI < 1.

Mehra suggested that individuals who are interested in short-term investment planning will wish to project the conditional equity premium over their planning horizon. But doing so is by no means a simple task. It is isomorphic to forecasting equity returns. Because returns have a standard deviation of 20 percent, the noise dominates the drift. Operationally, how much information comes from knowing that the mean risk premium is 2 percent rather than 6 percent when the standard deviation is 20 percent?

In conclusion, Mehra considered how the world must have looked to an investor at the end of 1928.

The mean real return on the S&P 500 had been 8.52 percent for 1889–1928, and the mean real return on risk-free assets had been 2.77 percent, so the observed mean equity risk premium would have been 5.75 percent (575 bps). An analysis similar to the Mehra–Prescott (1985) analysis, however, would have indicated an *ex ante* premium of 2.02 percent.

Is the future likely to be different from the past? To decide, we need to focus on what factors might make the future different. Demographic changes, for example, could be very important. But, maybe, because of market efficiency, the market has already taken into account the likely changes.

Current Estimates and Prospects for Change: Discussion

John Campbell (*Moderator*)
Ravi Bansal
Bradford Cornell
William Goetzmann
Roger Ibbotson
Martin Leibowitz
Rajnish Mehra
Thomas Philips
William Reichenstein, CFA
Stephen Ross
Robert Shiller
Jeremy Siegel

JOHN CAMPBELL (Moderator)

I'll make a few remarks and then open the discussion. I would like to amplify a distinction that Raj Mehra was making between the *ex post*, realized premium over some past period and the *ex ante* premium that investors are expecting at a single point in time. Over the long run, these premiums have to average out to the same level if the market has any rationality at all, but in the short run, they can move quite differently. For example, a lot of Raj's graphs indicate that the *ex post* and *ex ante* risk premiums might move in opposite directions, and I think that concept is very important to keep in mind. If we go through a period when the *ex ante* premium falls (for whatever reason), that movement will tend to drive prices up for a given cash flow expectation, so we will see a high realized return during a period when the *ex ante* premium has actually fallen. That is the story of the 1990s—that average returns were high, particularly at the end of the decade, because investors were willing to take on more risk, so the required rate of return was declining. Thus, we had a decline in the *ex ante* equity premium at the same moment that we had very high average returns.

Of course, if the equity premium is estimated by use of historical average returns, even over a period as long as 100 years, a few good years can drive up the long-term average considerably. For example, over 100 years, a single good yearly return of 20 percent adds 20 bps to the 100-year average return. This is the problem with estimating the equity premium from historical average returns; there is so much noise, and the average will tend to move in the wrong direction if the true *ex ante* premium is moving.

As a result, the methodology used by many at this forum is to focus on valuation ratios at a single point in time and make adjustments for growth forecasts. The methodology can be applied simply or elaborately. You can simply look at the earnings yield, or you can try to adjust the yield for return on equity being greater than the discount rate equilibrium or Tobin's q being different from 1, which we discussed this morning [in the "Historical Results" session]. I think this approach is the right way to go. If you want to estimate the *ex ante* premium, you start with a valuation ratio that summarizes the current state of the market, make some adjustments based on your best judgment, and back out the *ex ante* premium.

The approach has two difficulties that one has to confront. They arise from the fact that the models we are using are steady-state models that give long-term forecasts in a deterministic setting. The problem with using a deterministic model is that you obliterate any distinction between different kinds of averages. In a random world, however, that distinction matters a lot. It matters to the tune of 1.5–2.0 percentage points.

The second problem is that a forecast from a valuation ratio is really the equivalent of the yield on a long-term bond. The valuation ratio produces an infinite discounted value of future returns. You don't necessarily know the sequence of predicted returns. You don't know the sequence of forward rates or the term structure; you just have a single measure of a long-term yield. So, it's very difficult to construct or generate a view about the actual path that returns might follow.

In my work with Bob Shiller, we argue that, given the level of prices, this long-term yield must be very low. But that argument is consistent with two different views about the time path. One view is that a correction is going to occur in the short or medium term, followed by a return to historical norms. If you hold this view, you have to be bearish in the short term but you are more optimistic about returns in future years. This outlook would be very pessimistic for an investor who has finished accumulating wealth and wants to cash out; it would be a more optimistic outlook for an investor who expects to accumulate assets over the next several decades.

The other view, which I think has some plausibility, is that we might see mediocre returns over the long term because of structural changes—structural changes in that transaction costs have come down, the costs of diversification have come down, investors have learned about the equity premium puzzle, and therefore, the *ex ante* premium is down and will be permanently down. This view is less bearish in the short term than the first view but also less optimistic in the long term.

I think Bob and I differ a little bit on this time-path issue in terms of how to chop up the long-term yield into a sequence of forecasts. Bob is probably closer to the view that returns will be very poor in the short term and then revert to historical norms, and I am closer to the view that there may have been a permanent structural change that will mean mediocre returns in the near term and the longer term.

It is hard for me to imagine a long-run equilibrium with an equity premium relative to U.S. T-bills less than about 1.5 percent geometric (2.5–3.0 percent arithmetic). And I think it may take a further price decline to reach that long-run equilibrium. In other words, we are in for a short period of even lower returns followed by a (geometric) premium of about 1.5 percent for the long term.

MARTIN LEIBOWITZ: One thing we have not talked much about is that if, over time, we have more data on earnings, price movements, and returns, what is going to be the catalyst for moving the risk premium to higher or lower levels—or to a point of acceptance? Of course, one of the really great things about the market is its ambiguity; even if you are earning dismal returns now, the market's volatility always allows you to look back at a recent period when you earned great returns. But what sequence of events and flow of information would wake up market participants to say, "Hey, a 2 percent equity risk premium? I'm not buying for 2 percent. Give me something else. Is there another market I can invest in? Is there another advisor out there?" This possibility is worth thinking about because if we make the rounds and tell our friends and professional colleagues, "Look, we've found out that the nominal, arithmetic equity risk premium is roughly only 3.0–3.5 percent, and that's going to be it, but I can give you some good news: Volatility will be relatively low, so you will really be getting a lot of return for the amount of risk you'll be taking," people will say, "Forget it!" I would not want to be invested in the equity market with that sort of outlook. People would just run away from the equity market. People are thinking, hoping, and dreaming of returns well over an equity premium of 3 percent; they are thinking of a risk premium greater than that. This kind of question is what we need to discuss.

RAJNISH MEHRA: This point is the reason that understanding *why* we have an equity premium is so important. On the one hand, if there is a rational reason for the equity premium—for instance, if investors are scared of recessions and actually demand a 6 percent equity premium, then I would expect a 6 percent premium in the future. On the other hand, if we find out that investors do not actually demand that premium for holding stocks—that they perceive stocks, in some sense, to be not much riskier than bonds—then, the premium will be lower. You seem to be saying that investors *do* perceive stocks to be much riskier than bonds and they *do* want a high premium, in which case they will get it. If investors refuse to own stocks when they get only a 2 percent premium, a repricing of assets will take place.

STEPHEN ROSS: One thing that we all agree on is that there is enormous estimation error in figuring out the risk premium. I find it ironic that the estimation error in the risk premium that we agree on plays no role whatsoever in the models that we use to infer the risk premium. It is somewhat like option pricing, where you assume you know the volatility. You look at the option price, and then you figure out what the volatility must be for that to be the option price. Then, you build models of what the option price should be. But estimating the risk premium is even more complicated, and estimation error is even more damaging.

The estimation error in estimating the risk premium is huge. Over a 100-year period, the standard error alone of the sample estimates is on the order of 2–3 percent. I am not convinced by John Campbell's argument that structural models, which are efforts to get conditional probability estimates and do a better job of conditioning, will improve the situation, because we have about the same volatility on our conditional estimates. I have a very pessimistic view of those models. They introduce other parameters, and where we had 2 percent standard errors on a few parameters, now we have 4 percent because we have more parameters. I'm not convinced that this approach will narrow down the estimate.

I am troubled by the fact that in this world of incredible volatility, and with no real confidence in our estimations of the risk premium, we still go ahead and advise people about what to do with their portfolios. As Rajnish Mehra said, we have a strange disconnect: The uncertainty that we all perceive in these models plays no role in the construction of the models. As a consequence, uncertainty plays no role in our ability to filter from the models better estimates. One of the things we have to think seriously about is estimation error in these models.

THOMAS PHILIPS: I share John Campbell's view that, barring an unforeseen surge in productivity, we are in for a prolonged period of lower returns prior to transaction costs and fees. However, the actual return that will be *realized* by investors net of transaction costs and fees is probably not very different from the return achieved in the past. Don't forget that index funds did not exist in 1926. In those days, transaction costs and fees subtracted 2–3 percent each year from returns; today, costs have fallen by 90 percent.

WILLIAM REICHENSTEIN: A number of models predict returns using a dividend model. In this model, long-run return is the current dividend yield plus long-run expected growth in dividends plus the percentage change in price divided by the dividend multiple, P/D. When predicting returns, analysts tend to drop the last term and predict the capital gains as the long-run growth in dividends. In the corresponding earnings model, predicted return is the current dividend yield plus the capital gains (the long-run growth in earnings) plus the percentage change in P/E. That has to hold; it is a mathematical certainty.

The reason I do not like the dividend model but like the earnings model is that we have no idea where the P/D multiple is going to go. Yet, the predictions from the dividend model assume it will remain constant. I can accept that there is some normal range for the P/E multiple, but I agree with Fisher Black that there is no normal range for the P/D multiple. Black looked at the various arguments to try to explain why companies pay dividends, and in the end, he threw up his hands and said we have no idea. If we have no theory or empirical evidence to explain dividend policy, then we have no reason to believe the P/D multiple is going to be stable. And we have no way of predicting it. That ratio could go to infinity. Therefore, any model that drops out that term, even for a long-run analysis, may be very, very wrong.

BRADFORD CORNELL: The dividend ratio may not be stable. In fact, we are seeing declining dividends, but you may have a constant payout ratio.

REICHENSTEIN: If we wanted to estimate the ending P/E after the next 50 years, whatever we came up with, we might feel reasonably confident it is going to be between 30 and 8.

ROSS: It is higher than 30 now!

REICHENSTEIN: Let's say that something will stop the P/E multiple from going too high or too low. But if you ask what the ending P/D multiple will be, well, if companies keep dropping dividends, it could be a billion.

CORNELL: That is why you might want to include payouts. Wouldn't you think that political pressures would arise to make sure shareholders got a certain fraction, on average, of corporate earnings? If shareholders do not get some share, they will become dissatisfied and companies will not be able to issue equity. Corporations cannot play the game of siphoning off all the earnings indefinitely for executives' perks and options and so forth.

ROGER IBBOTSON: You do not have to get your return through dividends. If the company is bought out, you can get your money out. You can get your money out in lots of ways other than dividends. Speaking for myself, if I had a choice, I would not want to get any of my money out in dividends.

MEHRA: Tandy Corporation, for instance, does not pay out any dividends. It was sued by the U.S. IRS, which charged that it was helping stockholders evade taxes. The company successfully won the case with an argument that it had a diverse group of stockholders and was not acting in the interest of any particular shareholder group. A rational approach would be for shareholders, instead of receiving a dividend payment, to sell shares and pay a capital gains tax when they want cash.

REICHENSTEIN: Yes, we do end up paying taxes. So, if you are only able to tell me that 50 years from now, the P/D multiple could be anywhere from infinity to something much, much lower, then that is a heck of an estimation error.

ROSS: The interesting question being raised is whether price to dividends is the variable you should be looking at or whether we should be asking: Is there stability in price divided by total payout, including stock repurchases, dividends, and Roger Ibbotson's suggestion that there is a constant probability that you will get a cash offer for the holding? So, the totality of all the payouts would be an interesting long-term variable to look at that may well be quite stable.

CORNELL: There are also some monies that go the other way, however, so the effective payout rate is very hard to compute.

REICHENSTEIN: But if you are using a model and put in the current dividend yield to project long-run growth and if dividends come from some historical average, then in a period like the past 20 years (in which we have had this dramatic fall in dividend payout rates and dividend yields), if you don't include repurchases, you have a problem. Past growth is going to be below future growth, and the dividend model predictions miss this point. I think Stephen Ross is saying that dividend payouts are unstable but might be stable if we added back in repurchases. In my view, the dividend model is a questionable framework.

RAVI BANSAL: Both Rajnish Mehra and Bob Shiller commented on the size of the premium but didn't comment on, or make predictions about, the underlying volatility of the market portfolio. From John Campbell's comment, if I am interpreting it correctly, he views the current scenario as a form of a drop in the Sharpe ratio. Has uncertainty fallen or risen? What is happening to the Sharpe ratio?

CAMPBELL: There haven't been any long-term trends in the volatility of the market as a whole. Certainly, marketwide volatility fluctuates. Volatility was unusually low in the mid-1990s and has risen a lot since then, but if you look over decades, you don't see any trend. The result is different when you look at the idiosyncratic volatility measure, however, because then you do see a trend over the last three decades. But looking marketwide, we do not see trends. Actually this lack of trend is a puzzle because of the evidence that the real economy has stabilized. GDP growth seems to be less volatile. So, some people claim that risk has fallen, which would justify the fall in the equity premium. Yet, we don't see that lower volatility when we look at short-term stock returns. The market does not appear to think that the world is any less risky.

JEREMY SIEGEL: Could I suggest something? Because real uncertainty has declined, companies can lever up more, generate higher P/Es. The result is maintenance of equity volatility, but it's because of an endogenous response to the increased real stability of the economy. So, greater leverage and higher P/Es could be generating the same equity volatility, which wouldn't be a puzzle even with the more stable real economy.

CAMPBELL: But if companies have levered up to maintain the same equity volatility, the equity premium should not fall as a result.

SIEGEL: Yes, if you don't take labor income being more stable into account as one of the factors that might determine risk preferences. In fact, some research shows that if there were more stability on the wage side (labor income), that stability would give people more incentive to buy equities.

WILLIAM GOETZMANN: Just a word on dividends: With all the studies that have looked at historical dividend yields, the problem is that we do not know very much about the dividends on which the studies were based. For data before 1926, we have the Cowles Commission (1938) information on dividends, but when you start reading Cowles' footnotes, you see he had a problem figuring out whether he was actually identifying all the dividends that were being paid by the companies.

ROBERT SHILLER: Have you solved this problem? We had the same problem.

GOETZMANN: Well, no, but we found it was a striking problem. We started from the Cowles period and worked back to see if we could collect information on dividends. We have the information back to the 1820s or so, but we could be missing dividends.

SHILLER: You're concerned that you don't have all the information, that you are missing a significant chunk of it?

GOETZMANN: Yes. You have a set of stocks that are similar to each other—their industrial characteristics are similar, for example. One stock may be paying 8 percent dividends for 10 years, but for another stock, you have no dividend information available. Are you to presume that the second stock did not pay any dividends or that your records simply do not show the dividend? So, what we have had to resort to is to report the high number and to report the low number. And we don't think anybody else has ever really been able to get any better information about dividends than we have. So, if we're going to talk about model uncertainty, let's also talk about data uncertainty—particularly as the records go back through time.

SHILLER: Do you think that companies sometimes reported dividends to commercial and financial chronicles and at other times, misreported them or didn't report them at all?

GOETZMANN: Yes, that could be true.

SHILLER: Wouldn't it have to happen on a big scale to affect the aggregate numbers?

IBBOTSON: As you go back in time, it is not clear who or what was getting the reports. For one period of time, there was an official source for the NYSE, but later, that source disappeared. It is hard enough to get actual stock price data, but it is much harder to find out who reported dividends to whom. Therefore, dividend information comes from all sorts of sources.

GOETZMANN: So, for what it's worth, sprinkle some more noise into this whole process. It's a real challenge to focus on valuation ratio regressions. We've been talking about valuation ratio regressions and statistics in one form or another for eight or nine years now, and we have all sorts of details about the econometrics, but the real issue to me is whether we really know what the payouts were as we push backward in time.

IBBOTSON: For the stock price data, we only needed to go to one (or possibly two or three) sources, but for the dividend data, we had to go to many sources, and even after going to many sources, we found we were getting only some of the data. However, when we found the data, companies paid all their earnings out in dividends. They had 100 percent payout ratios in the 19th century. But for the missing data—who knows.

ROSS: In this entire discussion, we are focusing entirely on the risk premium, and we have sort of ignored the other variable, volatility. What is interesting about volatility is that it is the one variable about which we do have confident expectations.

Volatility has two features that are curious. One feature is that we can actually measure volatility with a certain amount of precision; we know what volatility is. Volatility is a lot less ambiguous than the equity risk premium. We need to bring volatility to bear on such questions as long-run portfolio allocation problems. Someone who has great estimation error about the risk premium and cannot quite figure out what it is but who, nonetheless, is taking others' advice as to what to do, would perhaps be informed in this decision by observing that we do know a lot about the pattern of volatility, we have far less estimation error for it, we sort of know what volatility is today, and we have pretty good ability to predict it over fairly long horizons. At least this person should understand the volatility of volatility, which shows up as much in those allocation problems as does expected return.

The second curious feature of volatility is, it seems to me, that we can use this variable in some interesting ways. Implied volatilities have been around now for 20 years. I know that the week before the 1987 crash, implied volatilities went to an annualized rate of about 120 percent. Prior to the current crash, implied volatilities again rose substantially. The cynic would say, well, implied volatility was quite high, but people didn't know whether the market was going up 200 points or down 200 points the next day; they just knew it was going to be a big move. But my guess is that investors figured that the market wasn't going to go up much more; they really thought the market was going to go down. It would be nice for those who are doing the empirical work on the risk premium to have a variable that actually has expectation recorded in it. It might be fun to look at its empirical content for the puzzles we are talking about today.

SIEGEL: I would like to add something to that comment. I think we know short-run volatility because we can measure it using options, most of which are very short term. But the question of long-run volatility depends very much on the degree of mean reversion, which is very important for long-term investors and is, as we all know, subject to great debate.

ROSS: Actually, I suspect long-term volatility is subject to less debate than long-run returns. For short-run volatility, even for an option one year out, with pretty good liquidity, you can start to see reversion—pretty clear reversion—one year out.

SIEGEL: But we don't have 10-year, or 20-year, or 30-year options, which might be very important for longer-term investors.

ROSS: Volatility is a lot better measure than returns, for which we have nothing that tells us anything about the short term or the long term.

SHILLER: I want to remind you of the very interesting discussion in Dick Thaler's talk this morning about *perceived* volatility [See the "Theoretical Foundations" session]. We seem to be forgetting about the distinction between the actual and the perceived risk premium. When Marty Leibowitz was saying that people would not be interested in stocks with an equity premium of 1.5 percent, he may have been assuming that the perceived volatility was very high. Dick was saying that it is the *presentation* to the general public that affects the public's perception of volatility. His research disclosed a very striking result, which is that when you present investors with high-frequency data, they have a much different perception of what the data are saying than when you present them with less-frequent—say, annual—data. And the way the data are being presented is changing. When I walk down the street now, I can look up at a bank sign that alternates between time, temperature, and the Nasdaq.

LEIBOWITZ: I have a couple of comments. First, if you had a volatility estimate that you could live with and you had actual asset allocations that were stable and common—most asset allocations, at least by institutional investors, are surprisingly stable and common—you could (theoretically) clearly back out from those variables the implied risk premium. No big challenge. At least, you could back out mean–variance estimates. Of course, the question is: What kind of time horizon would you be looking at? The horizon would be the critical ingredient. If you were looking over a long enough time horizon, the risk premium could be 0.1 percent. If you were looking over a short horizon, the risk premium could be something enormous.

Robert Merton wanted me to introduce along these lines the Zvi Bodie construct.[1] Bodie says that the kind of option you would have to buy as you go out to very long horizons is very different, in terms of the Sharpe ratio, from a short-horizon option; it is a very expensive option. That reality has to tell you something.

The other thing that I want to mention is that the issue of equilibrium payout ratios is very important. The question is: When an equilibrium is reached, at which point earnings are growing at either the growth rate of the economy or near that rate (i.e., that rate is your stable equilibrium view), then in terms of dividends, how much of a company's aggregate earnings have to be put back into the company to sustain that growth? This is the critical question. All else would then follow from the answer. It's surprising that this issue has not been much addressed, as far as I know, even from a macro level.

[1] Robert Merton was invited to attend but could not.

PHILIPS: There is a pragmatic solution to the question that Stephen Ross and Jeremy Siegel raised. We have about 20 years of option data, so you might construct the volatility data going back 20 years, and you could explore the fact that as you sample faster and faster, the estimates of volatility get sharper and sharper. Just take a perfect-foresight model: Assume it's 1920, and you're going to assume that the world is rational and that the forecasted volatility would have been the volatility that was actually realized over 1921, or 1921–1925, or whatever years you want to use. From those data, you could impute a data series going back in time and then try to do the appropriate tests. Cliff Asness has a very nice paper in the *Financial Analysts Journal* that explores this approach (2000b). Cliff looks at historical volatility and then backs out future returns as a function of historical volatility.

Implications for Asset Allocation, Portfolio Management, and Future Research I

Robert D. Arnott

First Quadrant, L.P.
Pasadena, California

A practitioner's empirical approach to estimating prospective (expected) equity risk premiums does not bode well for finding alpha through conventional U.S. equity allocations. In the United States and the United Kingdom, real earnings and real dividends have been growing materially slower than real GDP. Based on empirical evidence, if today's dividend yield is 1.7 percent and growth in real dividends is about 2.0 percent, cumulative real return on stocks will be about 3.7 percent. With a 3.4 percent real yield on bonds available, the *ex ante* risk premium all but disappears. Perhaps most troubling in the empirical evidence is the 60 percent negative correlation between payout ratios and subsequent 10-year earnings growth. With current payout ratios close to 40 percent, the implication for earnings growth over the coming decade is a rate of about –2 percent. When an assumed negative earnings growth rate is combined with an assumed zero risk premium, we have a serious problem.

I have to begin by offering profuse apologies. You are seasoned, very capable academics, and I'm not. I'm just a practitioner and an empiricist. So, we're going to focus on practice and empiricism in this presentation and stay far away from the theory related to the equity risk premium.

History versus Expectations

First, I want to emphasize an observation that a number of speakers have made: Much of the dialogue about the risk premium is very confused because the same term, "risk premium," is used for two radically different concepts. One is the historical excess return of stocks relative to bonds or cash, and the other is the prospective risk premium for stocks relative to bonds on an *ex ante* basis, without any assumptions about changes in valuation levels. The two concepts are totally different, should be treated separately, and, I think, should carry separate labels. Excess returns measure past return differences. The risk premium measures prospective return differences. I wish the industry would migrate to using different terms for these two radically different concepts.

A quick observation: If you are a bond investor and you see bond yields drop from 10 percent to 5 percent, and in that context, you have earned a 20 percent return, do you look at those numbers and say, "My expectation of 10 percent was too low. I have to ratchet my expectation higher. I'll expect 12–15 percent"? Of course not. The reaction by the bond investor is, "Thank you very much for my 20 percent returns; now, I'll reduce my expectation to 5 percent." If the earnings yield on stocks falls from 10 percent to 5 percent, however, what is the investment community's response when they see the 20 percent return? They say, "Our expectations were too low! Let's raise our expectations for the future."

My impression of the discussion we have been having today is that the reaction in this room would be absolutely unanimous in saying the portion of return attributable to the drop in the earnings yield (earnings to price) or the drop in the dividend yield can and should be backed out of the historical return in shaping expectations. I haven't heard a lot of discussion of the fact—and I think it is a fact—that a drop in the earnings yield should have a second-stage impact. The first stage is to say 10 percentage points (pps) of the return came from falling earnings yields; therefore, let's back that out. The second stage is that the fall in the earnings yield should produce a haircut in future expectational returns. I don't hear this concept out in the marketplace, and I don't hear it much in the academic community either.

Strategic Implications of Lower Returns

Let's begin with the hypothesis that the risk premium, the forward-looking premium, on U.S. stocks is now zero. Please accept that supposition for the next few minutes. If the risk premium is zero, what is the implication for asset allocation policy? In the past, the policy allocation to stocks and fixed income was the king of asset management decisions. It was the number one decision faced by any U.S. institutional investor—indeed, any investor in general. The reason was that more stocks meant more risk and more return.

The fiduciary's number one job was to gauge the risk tolerance of the investment committee and to push the portfolio as far into stocks as that risk tolerance would permit. If that job was done correctly, the fiduciaries had succeeded in their primary responsibility. But if stock, bond, and cash real returns are similar, if the risk premium is approximately zero, then it doesn't matter whether you have a 20/80 equity/debt or an 80/20 equity/debt allocation. It does affect your risk and your year-by-year returns, but it doesn't affect your long-term returns. So, if the risk premium is zero, this fundamental policy decision is radically less important than it has ever been in the past.

As for rebalancing, the empirical data support the notion that rebalancing can produce alpha, but we do not have a lot of empirical data to support the notion that rebalancing adds value. History suggests that rebalancing boosts risk-adjusted returns, but it sometimes costs money. Rebalancing produces alpha by reducing risk, and in the long term, it typically adds some value in addition to risk reduction. Now, suppose we are in a world in which there is no risk premium and in which stocks and bonds have their own cycles, their own random behavior. If that behavior contains any pattern of reversion to any sort of mean, rebalancing suddenly can become a source not only of alpha but also of actual added value—spendable added value.

In the past, tactical asset allocation (TAA) provided large alpha during periods of episodic high returns but did not necessarily provide large added value. So, the actual, live experience of TAA in the choppy, see-saw market of the 1970s was awesome. In the choppy bull market of the 1980s, value added from TAA was not awesome but was still impressive. In the relentless bull market of the 1990s, the value added from TAA was nonexistent. Alpha was certainly still earned in the 1990s (a fact overlooked by many), but it came mostly from reduced risk. If we are moving into markets like those of the 1970s, then TAA certainly merits another look.

What about the strategic implications of lower returns for pension funds? If conventional returns lag actuarial returns, then funding ratios are not what they seem. I did a simple analysis of funding ratios for the Russell 3000 Index and found that for every 1 pp by which long-term returns fall short relative to actuarial returns, the true earnings of U.S. pension assets fall by $20 billion. If, as I believe is the case, long-term returns are going to be about 3 pps below long-term actuarial assumptions, pension fund earnings will be $60 billion less than what is being reported, and this shortfall will need to be made up at some later date.

In a world of lower returns, if you don't believe in efficient markets, alpha matters more than ever before. If you do believe in efficient markets, the avoidance of negative alpha by not playing the active management game matters more than ever.

Now, a truism would be that conventional portfolios will produce conventional returns. That is fine if conventional returns are 15 percent a year, as they were for the 18 years through 1999. In a market environment of 15 percent annual returns, another 1 pp in the quest for alpha doesn't matter that much to the board of directors, although it does make a material difference to the health of the fund. However, if the market environment is producing only 3–4 percent real returns for stocks *and* bonds, another 1 pp matters a lot.

What investments would be expected to consistently add value in a world of lower expected returns? "Conventional" alternative investments may or may not produce added value. Private equity and venture capital rely on a healthy equity market for exit strategies. They need a healthy equity market to issue their IPOs (initial public offerings). Without a healthy equity market, private equity and venture capital are merely high-beta equity portfolios that can suffer seriously in the event of any sort of reversion to the historical risk premium. International equities and bonds may have slightly better prospects than U.S. equities and bonds, but not much better.

Strategies well worth a look are the elimination of slippage, through the use of passive or tactical rebalancing, and cash equitization. If the equity risk premium is lost, then alternative assets whose returns are uncorrelated with the U.S. equity market will absolutely produce added value. Uncorrelated alternatives include TIPS,[1] real estate, REITs (real estate investment trusts), natural resources, and commodities. Absolute return strategies (market-neutral or long–short strategies and other hedge fund strategies) will also absolutely produce added value—if you can identify strategies that *ex ante* have an expectation of alpha. These approaches are, more than anything else, bets on skill and bets on inefficient markets. So, the investment strategies that will work in a world of lower returns differ greatly from the conventions that are driving most institutional investing today.

These reflections are from the vantage point of a practitioner. Much of what I've said makes the tacit assumption that markets are quite meaningfully inefficient, so these comments might be viewed with a jaundiced eye by a group that accepts market efficiency. Now, let's turn from practice to empiricism.

Empirical Experience

The Ibbotson data going back 75 years show about an 8 percent cumulative real return for stocks (see Ibbotson Associates 2001). Starting at the end of 1925 with a 5.4 percent dividend yield, the valuation attached to each dollar of dividends quadrupled in the 75-year span. That increase translated into nearly a 2 percent a year increase in the price/dividend valuation multiple—hence, 2 pp of the 8 percent real return. I think nearly everyone in this room would feel comfortable backing this number out of the returns in shaping expectations for the future. Over the 75-year period, real dividends grew at a rate of 1 percent a year. So, over the past 75 years, stocks produced an 8.1 percent real return. The real yield at the start of this period was 3.7 percent. (I say "real" yield because the United States was still on a gold standard in 1925; inflation expectations were thus zero. Bonds yielded 3.7 percent, and bond investors expected to earn that 3.7 percent in real terms.) Bonds depreciated as structural inflation came onto the scene. So, stocks earned a cumulative 4.7 percent real return in excess of the real return earned by bonds over the same period.

What does the future have in store for us from our vantage point now in the fall of 2001? **Table 1** contains the Ibbotson data and our analysis of the prospects from October 2001 forward. We'll start with a simple model to calculate real returns for stocks:

Real stock return = Dividend yield
+ Dividend growth
+ Changes in valuation levels.

In October 2001, the dividend yield is roughly 1.7 percent. If we assume that stock buybacks accelerate the past growth in real dividends, we can double the annual growth rate in real dividends observed over the past 75 years to 2 percent. Those two variables give us a 3.7 percent expected annual real return. TIPS are currently producing a 3.4 percent annual real return. Thus, the expected risk premium is, in this analysis, 0.3 pp, plus or minus an unspecified uncertainty, which I would argue is meaningful but not huge.

Why was the historical growth in real dividends (from 1926 through 2000) only 1 percent a year? Did dividends play less of a role in the economy? Were corporate managers incapable of building their companies in line with the economy? I don't believe either was the reason. The explanation hinges on the role of entrepreneurial capitalism as a diluting force in the growth of the underlying engines for valuation—that is, earnings and dividends of existing enterprises. The growth of the economy consists of growth in existing enterprises and the creation of new enterprises. A dollar invested in the former is not invested in the latter. **Figure 1** shows real GDP growth, real earnings per share (EPS) growth, and real dividends per share (DPS) growth since January 1970. Over the past 30 years, until the recent earnings downturn, real earnings have almost kept pace with real GDP

Table 1. The Ibbotson Data Revisited and Prospects for the Future

Component	75 Years Starting December 1925	Prospects from October 2001
Starting dividend yield	5.4%	1.7%
Growth in real dividends	1.0	2.0
Change in valuation levels[a]	1.7	???
Cumulative real return	8.1	±3.7
Less starting bond real yield	3.7[c]	3.4[d]
Less bond valuation change[b]	–0.4	???
Cumulative risk premium	4.7	±0.3

[a] Yields went from 5.4 percent to 1.4 percent, representing a 2.1 percent increase in the price/dividend valuation level.
[b] Bond yields went from 3.7 percent to 5.5 percent, representing a 0.3 percent annualized drop in long bond prices.
[c] A 3.7 percent yield, less an assumed 1926 inflation expectation of zero.
[d] The yield on U.S. government inflation-indexed bonds.
Source: Based on Ibbotson Associates (2001) data.

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

Figure 1. GDP, EPS, and DPS: United States, January 1970–January 2001

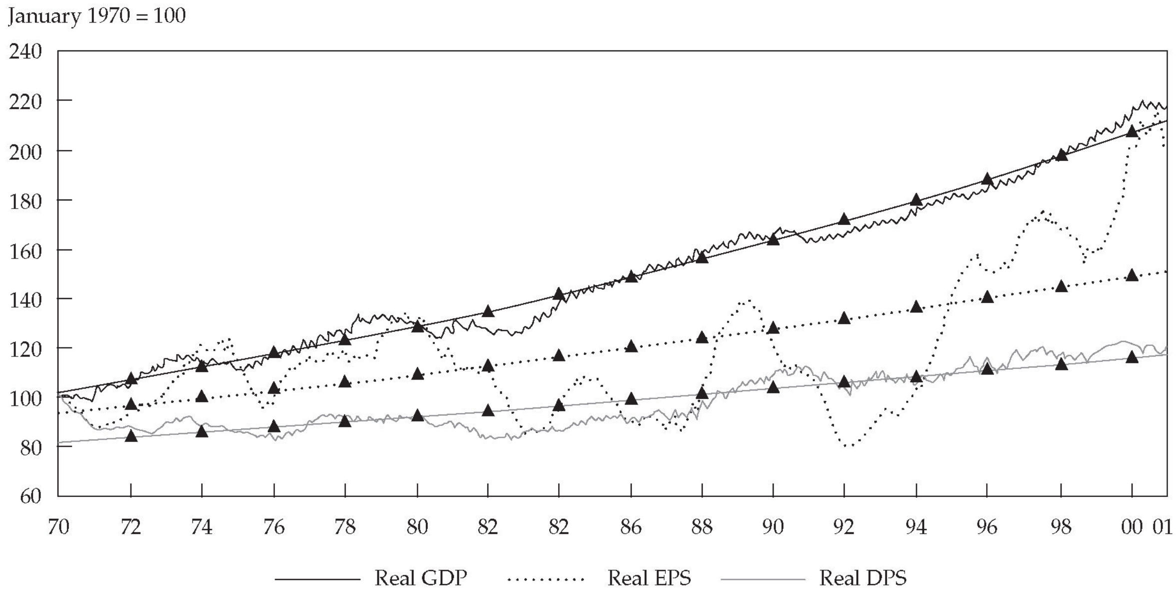

Note: Triangles identify exponentially fitted lines.
Source: Data from Organization for Economic Cooperation and Development (OECD).

growth. However, this pattern has occurred in the context of earnings as a share of the macroeconomy rising from below historical norms to above historical norms, including a huge boom in the 1990s. From the line of best fit, we can see that the growth trend in real earnings and real dividends is materially slower than the growth in the economy.

Is the picture different in Canada? Yes, it is. **Figure 2** illustrates that real earnings and real dividends on an indexed portfolio of Canadian equities have actually shrunk while real GDP has grown, producing a bigger gap between the series than we find in the United States. Why did this happen? In Canada, the fundamental nature of the economy has evolved in the past 30 years from resource driven to information and services driven.

The experience of the United Kingdom, where real earnings and real dividends grew materially slower than real GDP, has been similar to that of the United States. The experience of Japan has been rather more like Canada's. Japan, like Canada, is a fundamentally restructured economy. The result is that over the past 30 years, entrepreneurial capitalism in Japan has had a larger dilutive effect on shareholders in existing enterprises than it has in the United States.

Table 2 shows, for the period from 1970 through 2000, the average growth of the four countries in real GDP, real EPS, real DPS, and average real EPS plus real DPS; Table 2 also shows the combined averages for each country and for all four countries grouped together. The general pattern is clear: Entrepreneurial capitalism is the dominant source of GDP growth, so it dilutes the growth of earnings for investors in existing enterprises.

We can look back over a much longer span for the U.S. market, from 1802 to 2001. **Figure 3** graphs the growth of $100 invested in U.S. stocks at the beginning of the 200-year period. Assuming dividends are reinvested, the $100 would have grown to more than $600 million by December 2001—a nice appreciation in any portfolio. By removing the effects of inflation and reinvestment of dividends, we can isolate the internal growth delivered by the existing companies. When the effect of inflation is removed, the ending value drops to $30 million. And when the assumption of reinvested dividends is removed, the ending value is reduced to a mere $2,000.

Figure 4 illustrates the link between real growth in stock value and economic growth. Real GDP growth increased 1,000-fold over the 1802–2001 period, real stock prices increased some 20-fold, and real per capita GDP growth similarly increased about 20-fold.

We can now assess the underlying engines of valuation. We'll examine the real dividend (you could do the same thing with real earnings). As **Figure 5**

Figure 2. GDP, EPS, and DPS: Canada, January 1970–January 2001

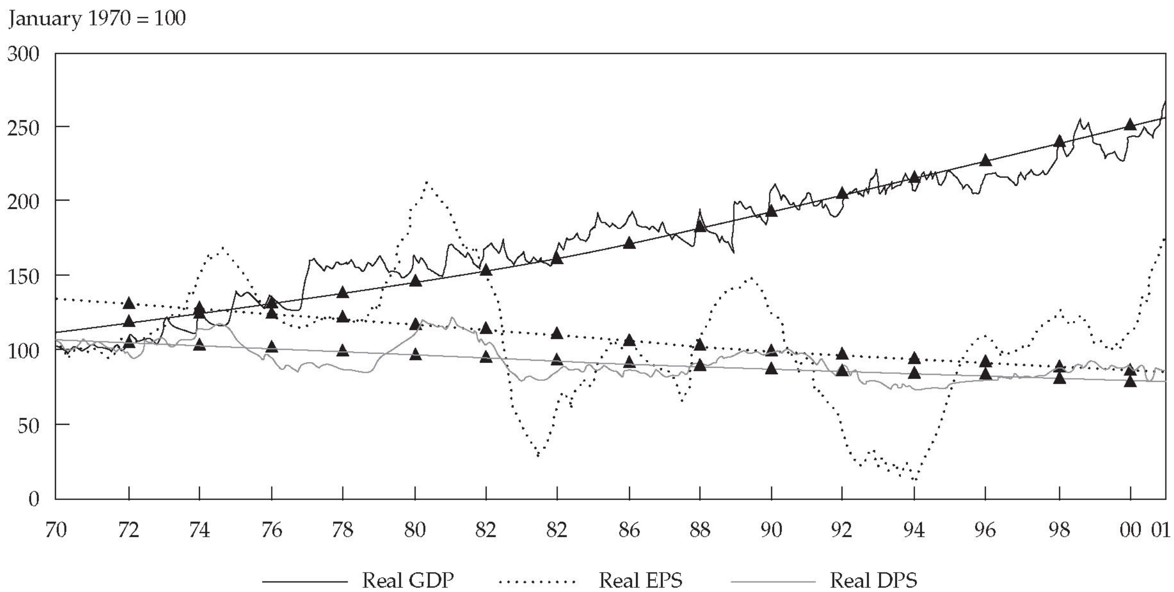

Note: Triangles identify exponentially fitted lines.
Source: OECD.

Table 2. Growth in GDP, EPS, DPS, and EPS + DPS, January 1970–January 2001

Measure	Canada	Japan	United Kingdom	United States	Average
Real GDP	2.7%	3.1%	2.4%	2.0%	2.5%
Real EPS	–1.4	–3.8	1.3	1.3	–0.6
Real DPS	–0.8	–1.6	2.0	1.0	0.1
Average real EPS + real DPS	–1.1	–2.7	1.6	1.1	–0.3
Average EPS + DPS growth as a percentage of GDP	–41.0	–87.0	67.0	57.0	–11.0

Source: OECD; Morgan Stanley Capital International.

shows, real dividend growth matches very closely the growth in real per capita GDP. The implication is that the internal growth of a company is largely a matter of productivity growth in the economy and is, in fact, far slower than the conventional view—that dividends grow at the same rate as GDP.

Now we are ready to model and estimate real stock returns. In **Figure 6**, the dashed line represents the dilution of GDP growth in the growth of dividends. Growth in dividends tracks growth in real per capita GDP (the dotted line) remarkably tightly; the standard deviation is very modest—only 0.5 percent. This relationship is astonishingly stable. On a 40-year basis, the deviation is never above +0.1 percent and never below –1.6 percent. Moreover, current experience is in line with historical norms, despite anecdotal opinions that companies are delivering less in dividends than ever before.

A model that estimates real stock returns is useful only if its estimates actually fit subsequent experience. **Figure 7** is a scattergram providing the correlation between estimated and subsequent actual 10-year real stock returns. The correlation between the two is approximately 0.46 for the full period and far higher since World War II. The current figure for the real stock return is down in the 2–4 percent range. Of course, what the subsequent actual real return will be is anybody's guess, but I am not optimistic.

The same type of modeling can be done to estimate the real bond return. An inflation estimate can be subtracted from the nominal bond yield to arrive at an estimated real bond return. How do the

Figure 3. Return from Inflation and Dividends, 1802–2001

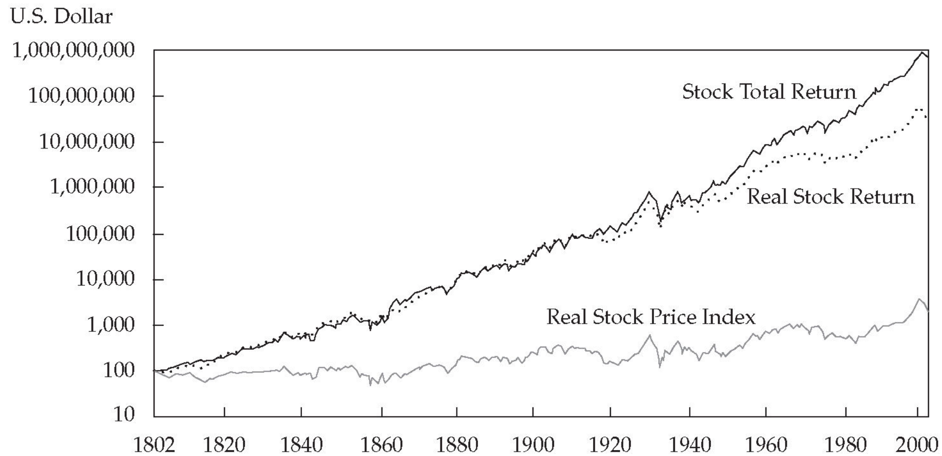

Notes: The "Real Stock Price Index" is the internal growth of real dividends—that is, the growth that an index fund would expect to see in its own real dividends in the absence of additional investments, such as reinvestment of dividends.

Source: Arnott and Bernstein (2002).

Figure 4. The Link between Stock Prices and Economic Growth, 1802–2001

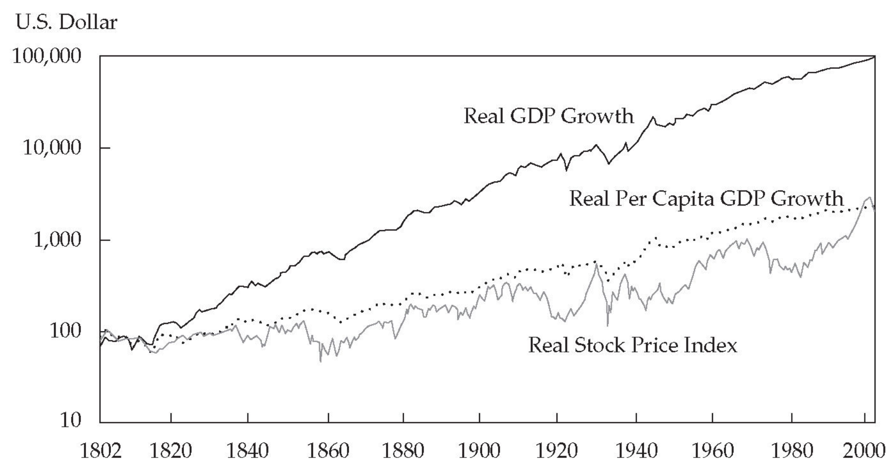

Source: Arnott and Bernstein (2002).

estimates calculated by this model fit with the subsequent real bond returns? As **Figure 8** shows, over a 200-year span, they fit pretty darned well. The loops off to the left relate to wartime. In several periods—the Civil War, World War I, World War II—investors were content to receive a negative expected real return for bonds, which can perhaps be attributed to patriotism. The country survived, so the real returns exceeded the expectations.

By taking the difference between the estimated real stock return and the estimated real bond yield, you get an objective estimate of what the forward-looking equity risk premium might have been for investors who chose to go through this sort of straightforward analysis at the various historical points in time. As shown in **Figure 9**, the *ex ante* risk premium of 5 percent, considered normal by many in the investment business, actually appears only during major wars, the Great Depression, and their aftermaths.

How good is the fit between this estimated risk premium and subsequent 10-year excess returns of

Figure 5. Dividends and Economic Growth, 1802–2001

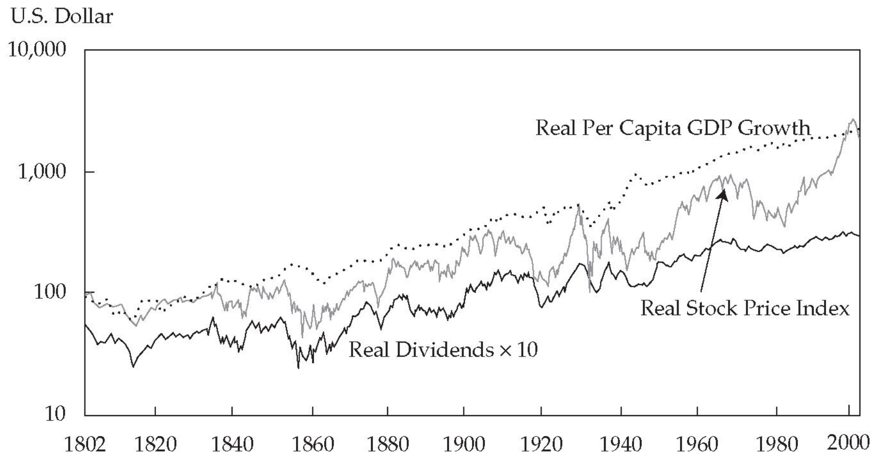

Notes: Real dividends were multiplied by 10 to bring the line visually closer to the others; the result is that on those few occasions when the price line and dividend line touch, the dividend yield is 10 percent.
Source: Arnott and Bernstein (2002).

Figure 6. Estimating Real Stock Returns, 1810–2001

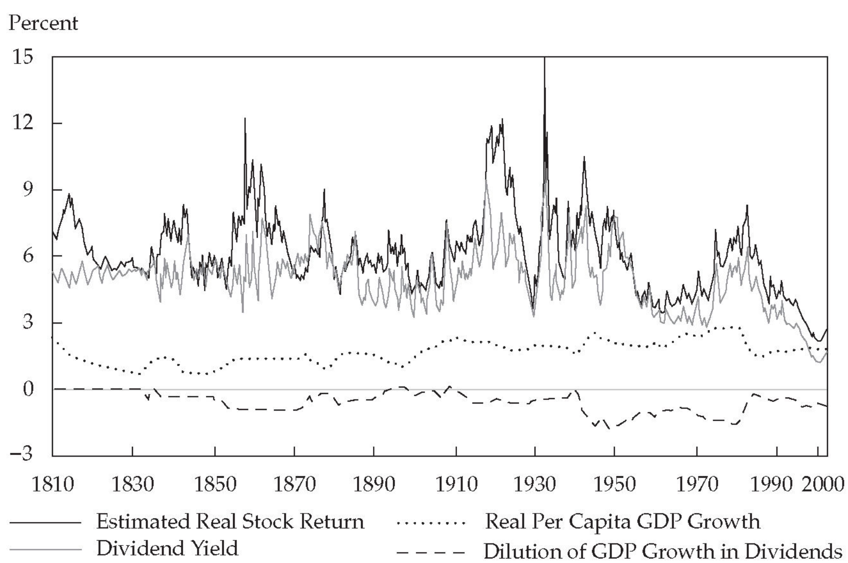

Notes: Based on rolling 40-year numbers. Real stock return = Dividend yield + Per capita GDP growth – Dividend/GDP dilution. The line "Dilution of GDP Growth in Dividends" indicates how much less rapidly dividends (and earnings) on existing enterprises can grow than the economy at large.
Source: Arnott and Bernstein (2002).

stocks over bonds? **Figure 10** shows that the fit is fairly good, which is worrisome in light of the poor current outlook. The current point on the *x*-axis (when this particular formulation is used) is about –0.5 percent. The implications for forward-looking 10-year real excess returns of stocks relative to bonds are worrisome—if this model holds in the future, if things are not truly different this time.

Figure 11 is a scattergram that relates the payout ratio to subsequent 10-year earnings growth from 1950 through 1991. This information ties in with Cliff Asness's talk [in the "Theoretical Foundations"

Figure 7. Estimated and Subsequent Actual Real Stock Returns, 1802–2001

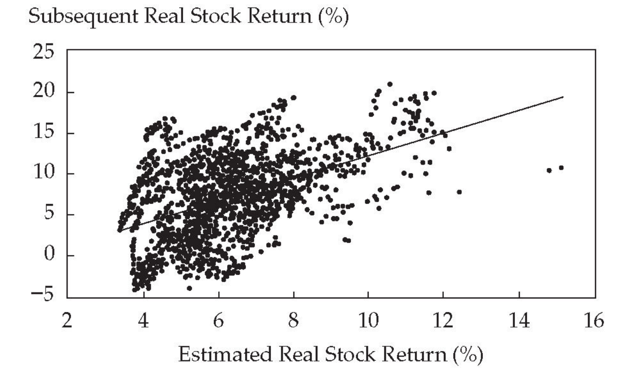

Source: Arnott and Bernstein (2002).

Figure 8. Estimated and Subsequent Actual Real Bond Yields, 1802–2001

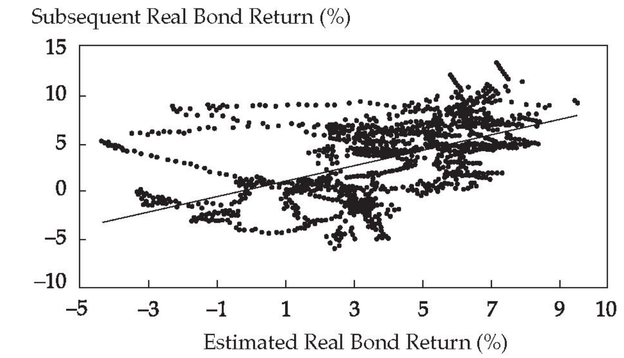

Source: Arnott and Bernstein (2002).

session]. Modigliani and Miller would suggest that if payout ratios are low (see Modigliani and Miller 1958), the reinvestment averaged across the market should produce the same market return that one could get by receiving those dividends and reinvesting them in the market. The tangible evidence is not encouraging. (Keep in mind that the M&M focus is cross-sectional, not intertemporal, so what I've just said is a variant of Modigliani and Miller's work, but it is a widely cited variant. M&M's work is frequently referred to in making the case that earnings growth is going to be faster than ever before.) Based on Figure 11, the correlation between payout ratios and subsequent 10-year earnings growth is a *negative* 0.60—which is worrisome. With recent payout ratios well below 40 percent, the implication for earnings growth is a rate of about –2 percent or worse, from the 2000 earnings peak, over the coming decade. If we combine an assumed negative earnings growth rate with an assumed zero risk premium, I believe that we have a serious problem.

Figure 9. Estimating the Equity Risk Premium, 1810–2001

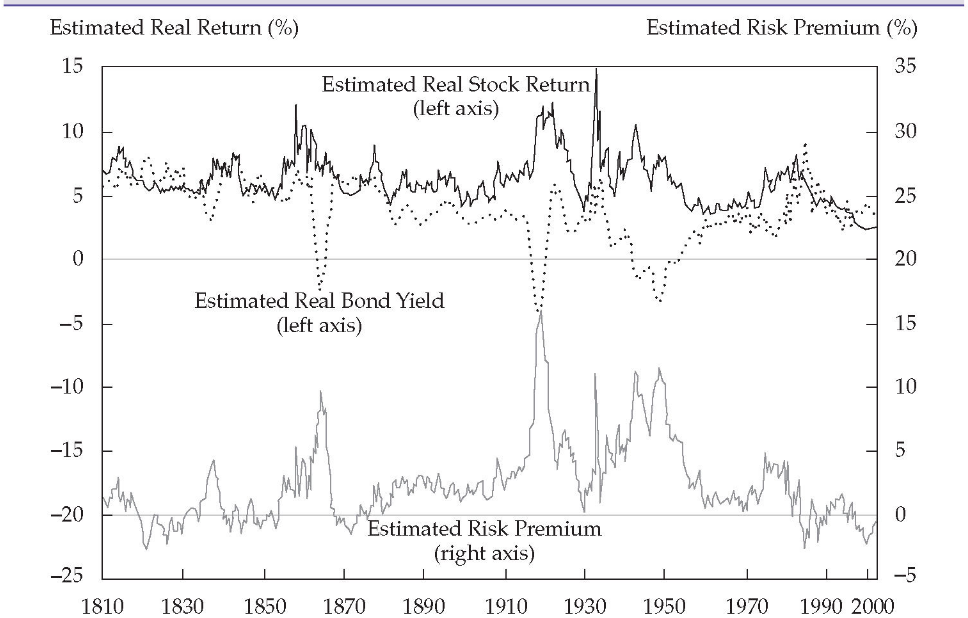

Source: Arnott and Bernstein (2002).

Figure 10. Risk Premium and Subsequent 10-Year Excess Stock Returns: Correlations, 1810–1991

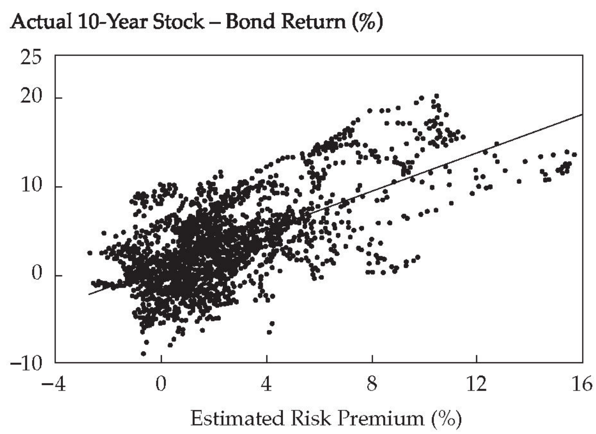

Source: Arnott and Bernstein (2002).

Figure 11. Payout Ratio and Subsequent 10-Year Earnings Growth, 1950–91

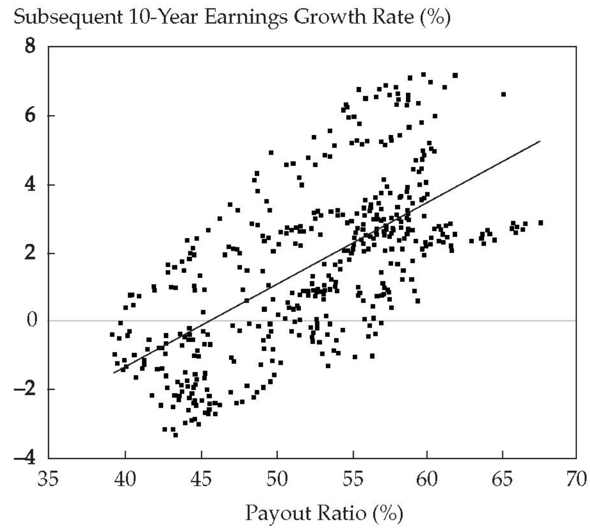

Implications for Asset Allocation, Portfolio Management, and Future Research I

Robert D. Arnott
First Quadrant, L.P.
Pasadena, California

SUMMARY

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

Robert Arnott began with an emphasis on practice and empiricism, as opposed to theory. He urged the use of the terms "equity excess return" for the past and "equity risk premium" for the future.

We have seen a decline in bond yields. Does this decline portend an increase or a decrease in bond returns? And we have seen a decline in stock earnings yields (earnings to price). Does this decline portend an increase or decrease in stock returns? The participants in the Equity Risk Premium Forum would all, he believes, when shaping expectations, back out the portion of return attributable to the drop in earnings or dividend yield from the historical return. But he had not heard much discussion of the fact that a drop in earnings yield should have a second-stage impact—a haircut in expected returns accompanying the fall in earnings yield.

Arnott estimated an *ex ante* risk premium at the present time of zero. In this case, the old policy of balancing risk and return no longer works. Rebalancing used to recognize that more stock meant more risk and more return. So, fiduciaries gauged the risk tolerance of the investment committee and pushed the portfolio as far into stocks as that risk tolerance would permit. If the return expectations for stocks and bonds are similar, the policy asset allocation matters in terms of risk but not in terms of returns and the allocation decision is far less critical than it was in the past.

Strategic Implications

Historically, rebalancing has produced an alpha by reducing risk. Over long periods, it produced a little extra return. Now, with no risk premium, with any pattern of reversion to a mean for stocks and for bonds, rebalancing can boost returns.

Tactical asset allocation achieved episodic returns that conveyed a large alpha in the turbulent 1970s and 1980s but did not necessarily add value in the roaring bull market of the 1990s, although it could reduce risk. If the U.S. market is headed for a repeat of the 1970s, then TAA may be especially worthwhile in the near future.

What about strategic implications for pension funds? If conventional returns lag actuarial estimates, which is likely, then current funding ratios are misleading, contributions will have to catch up, and alpha matters. In a world of lower returns, an emphasis on such alternative investments as private equity may be appealing, but to the extent that this emphasis relies on a strong equity market for an exit strategy, it may not be so attractive. International stocks and bonds may be attractive, but the expected returns there will also be low. Rebalancing and cash equitization are worth a look. Uncorrelated alternatives such as TIPS, real estate, REITs (real estate investment trusts), and commodities will be promising.[1] Absolute return strategies may be seen as more important in inefficient markets. There will be increased searching for inefficiencies by active managers and increased searching for avoidance of negative alpha by those who believe in market efficiency.

Empirical Results

Turning from practice to empiricism, Arnott's **Table 1** showed the Ibbotson data together with the prospects based on our current situation. Starting with a dividend yield of 5.4 percent, the U.S. equity market has seen an approximately 8 percent compounded real return on stocks over the past 75 years. The change in the price/dividend valuation ratio added 1.7 percent, which should be backed out of the returns for forecasting purposes. Note that real dividends grew at a scant 1 percent. The initial real bond yield in 1925 was 3.7 percent, and because it

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

Table 1. The Ibbotson Data Revisited and Prospects for the Future

Component	75 Years Starting December 1925	Prospects from October 2001
Starting dividend yield	5.4%	1.7%
Growth in real dividends	1.0	2.0
Change in valuation levels[a]	1.7	???
Cumulative real return	8.1	±3.7
Less starting bond real yield	3.7[c]	3.4[d]
Less bond valuation change[b]	–0.4	???
Cumulative risk premium	4.7	±0.3

[a] Yields went from 5.4 percent to 1.4 percent, representing a 2.1 percent increase in the price/dividend valuation level.

[b] Bond yields went from 3.7 percent to 5.5 percent, representing a 0.3 percent annualized drop in long bond prices.

[c] A 3.7 percent yield, less an assumed 1926 inflation expectation of zero.

[d] The yield on U.S. government inflation-indexed bonds.

Source: Based on Ibbotson Associates (2001) data.

was the quoted bond yield, investors had no reason to expect that inflation would matter. So, the excess return of equities over bonds was close to 5 percent. Now, we are looking at a 1.7 percent starting dividend yield, roughly a 2 percent growth in real dividends, and probably no increase in valuation levels—for a total prospective real return of about 3.7 percent. Subtracting a 3.4 percent real bond yield (e.g., the TIPS yield) produces a 0.3 percent (30 bps) cumulative risk premium plus or minus some small standard deviation.

Why did dividends grow at only 1 percent in the past? Looking at the **Figure 1** graph of real GDP, real EPS, and real dividends per share (DPS), we can see that earnings have almost kept pace with GDP growth—but in the context of going from a small share of the national economy to a large share. Entrepreneurial capitalism dilutes the growth experienced by investors in existing enterprises. The trend in dividend growth is well below that of GDP. Over the period January 1970 to January 2001, real GDP growth was fairly steady. Real earnings growth and real dividend growth followed slower trends and were quite irregular, with relatively high earnings growth since about 1995. The relative growth in GDP, equity earnings, and dividends has been similar in the United Kingdom to that in the United States. In Canada and Japan, however, the trend in earnings and dividends has been down, not up, over the past 30 years.

Turning to the 200-year history beginning in 1802, Arnott's **Figure 3** indicated that $100 invested in stocks in 1802 would have grown, with dividends reinvested, to nearly $1 billion in 200 years.[2] In real

[2] Table and figure numbers in each Summary correspond to the table and figure numbers in the full presentation.

Figure 1. GDP, EPS, and DPS: United States, January 1970–January 2001

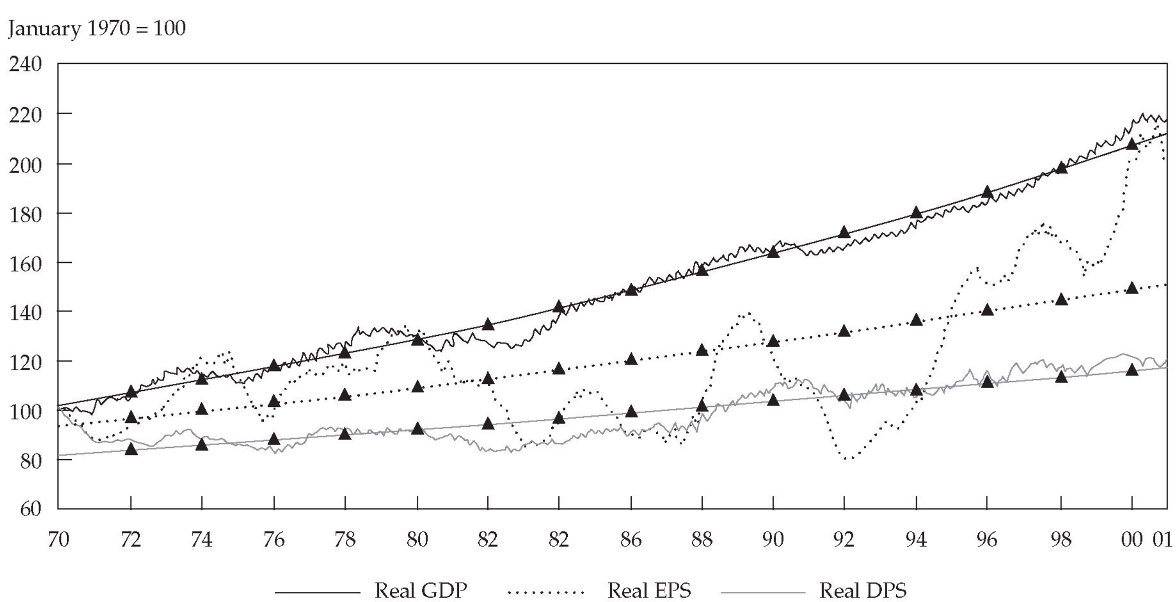

Note: Triangles identify exponentially fitted lines.

Source: Data from Organization for Economic Cooperation and Development (OECD).

Figure 3. Return from Inflation and Dividends, 1802–2001

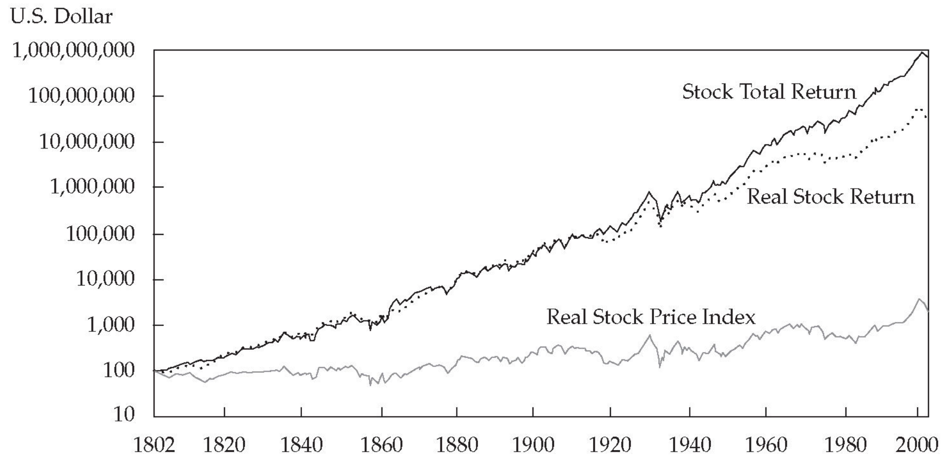

Notes: The "Real Stock Price Index" is the internal growth of real dividends—that is, the growth that an index fund would expect to see in its own real dividends in the absence of additional investments, such as reinvestment of dividends.

Source: Arnott and Bernstein (2002).

terms, however, the ending amount is $30 million, and when we look at the index alone, without dividend reinvestment, the $100 rose barely above $1,000.

Real dividends have trailed per capita GDP growth. **Figure 4** indicated that, in this time frame, an index of real stock prices tracked real per capita GDP growth rather well in the United States, although the index persistently trailed aggregate GDP growth for the 200 years.

Figure 6 provided a basis for modeling and estimating real stock returns. Real per capita GDP growth and dilution of GDP growth in dividends are both remarkably stable and closely parallel. The note to Figure 6 provides Arnott's equation for estimating real stock returns. This equation can also be used for the more recent subperiod of 1950–2001 to forecast future real stock returns. A similarly simple model can be used to estimate future real bond returns.

Figure 9 showed the results of using these simple models to estimate the real stock return, real bond yield, and equity risk premium (what might be called the "objective risk premium") year-by-year from 1810 to 2001. The risk premium rarely rose above 5 percent, only at the times of the Civil War, World War I, the Great Depression, and World War II. The premium is currently at or below zero.

During previous discussion of the Miller and Modigliani propositions, Arnott had commented that empirical evidence was not consistent with M&M. In this presentation, he showed the **Figure 11** plot of the payout ratio against subsequent 10-year earnings growth. Noting that M&M dealt with cross-sectional, not time-series, propositions and that he was showing time-series evidence, Arnott pointed out that high earnings retention (low payout) led *not* to higher earnings growth but to *lower* growth, a source of some concern.

Summary Implications

The implications of lower expected returns for policy allocation are as follows: In the past, the choice between stocks and fixed income was the essence of the policy asset-allocation decision. More stocks meant more risk and more return. For the future, with prospective stock and bond returns similar, policy allocation is no longer "king." If real earnings fall, as the empirical evidence on payout ratios suggests, or if valuation ratios "revert to the mean," then the situation is even worse.

Figure 4. The Link between Stock Prices and Economic Growth, 1802–2001

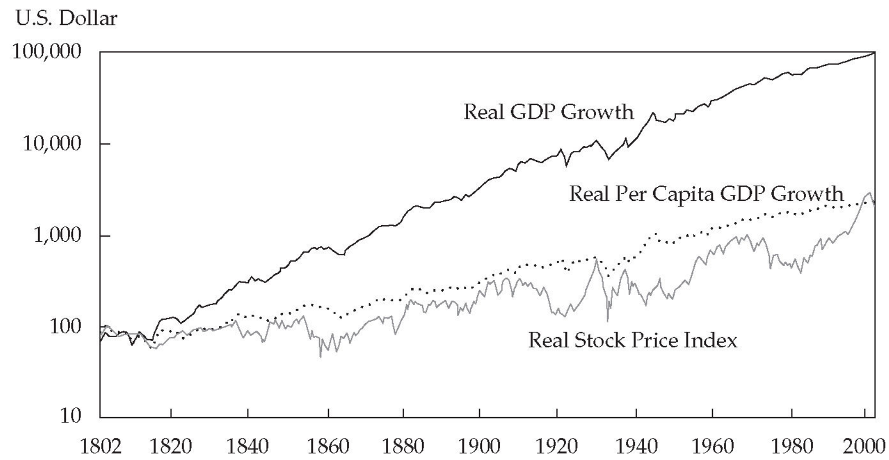

Source: Arnott and Bernstein (2002).

Figure 6. Estimating Real Stock Returns, 1810–2001

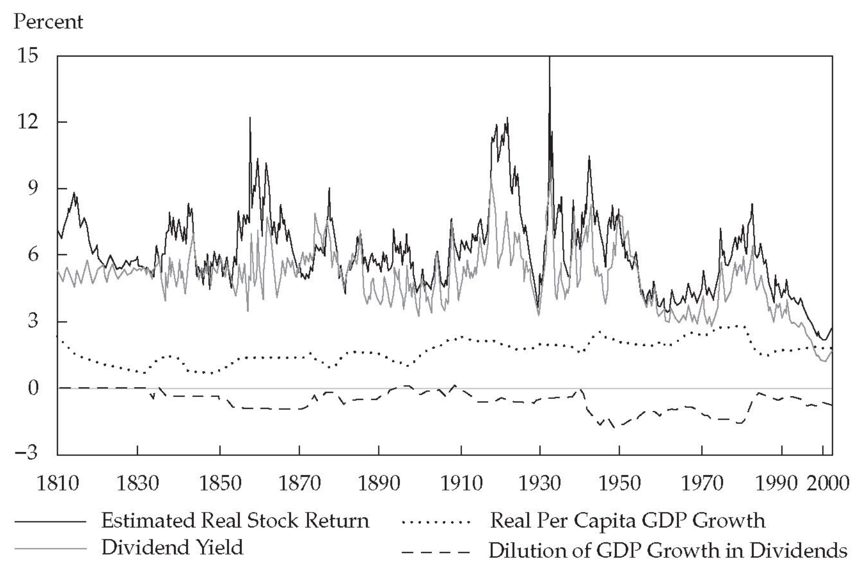

Notes: Based on rolling 40-year numbers. Real stock return = Dividend yield + Per capita GDP growth – Dividend/GDP dilution. The line "Dilution of GDP Growth in Dividends" indicates how much less rapidly dividends (and earnings) on existing enterprises can grow than the economy at large.

Source: Arnott and Bernstein (2002).

Figure 9. Estimating the Equity Risk Premium, 1810–2001

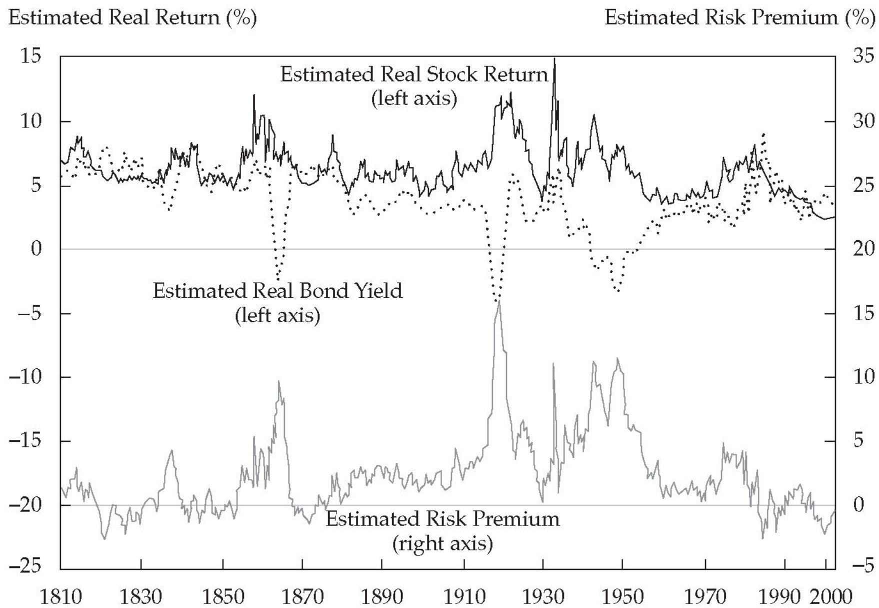

Source: Arnott and Bernstein (2002).

Figure 11. Payout Ratio and Subsequent 10-Year Earnings Growth, 1950–91

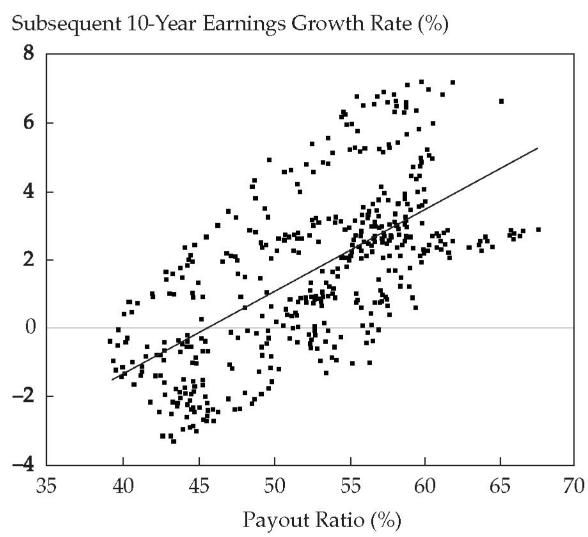

Implications for Asset Allocation, Portfolio Management, and Future Research II

Campbell Harvey

Duke University, Durham, North Carolina
National Bureau of Economic Research, Cambridge, Massachusetts

The reported survey of chief financial officers of U.S. corporations makes a unique contribution to the measurement of the expected equity risk premium and market volatility. Beginning with the second quarter of 2000, the research team has been conducting an ongoing, multiperiod survey of CFOs about their estimates of future equity risk premiums and equity market volatility. Results of the survey indicate the following: Return forecasts are positively influenced by past returns, which constitutes a type of "expectational momentum"; expected volatility is negatively related to past returns; the respondents seem to be very confident in their forecasts; and time horizon makes a big difference, in that a positive relationship was found between risk and expected return only for long-horizon forecasts.

After everything that has been said today, it is a challenge to make a unique contribution. We have heard how difficult it is to get a measure of expectations in terms of the equity risk premium, and what I am going to present is an approach to measuring expectations that is different from those that have been discussed.

For the past five years, John Graham and I, in conjunction with Financial Executives International, have been conducting a survey of chief financial officers of U.S. corporations about their estimates of future equity risk premiums and volatility.[1] Beginning in the second quarter of 2000 and, so far, extending into the third quarter of 2001, we have analyzed the more than 1,200 responses from the CFOs. Only 6 observations will appear in the graphs, but each observation is based on approximately 200 observations.

We know from other surveys that have been done that CFOs do actually think about the risk premium problem. We know that 75 percent of corporate financial executives—treasurers and CFOs—admit to using a CAPM-like or multifactor model. Therefore, we believe that the CFOs we are surveying are a reasonable sample of the population to question about the equity risk premium. I believe it is a sample group superior to that of economists surveyed—for example, by the Federal Reserve Bank of Philadelphia. The Philadelphia Fed's survey contains unreliable data (which I know from directly examining these data). I also think our survey has advantages over the survey of financial economists reported by Ivo Welch (2000) because our respondents are making real investment decisions. Finally, it is well known that the forecasts by financial analysts are biased. So, the survey we are conducting should provide some benefit in our search for *ex ante* risk premiums.

Survey of CFOs

Our survey has a number of components; it does not simply ask what the respondent thinks the risk premium is today. First, our survey is a multiperiod survey that shows us how the expectations of the risk premium change through time. Second, we ask about forecasts of the risk premium over different horizons. We have not talked much today about the effect of the investment horizon on the expected risk premium, but in our survey, we are asking about risk premium expectations for a 1-year horizon and a 10-year horizon. A third piece of information that we get in the survey is a measure of expected market volatility. Finally, we can recover from the responses a measure of the asymmetry or skewness in the distribution of the risk premium estimates.

[1]For a complete description of the study reported here, see Graham and Harvey (2001a).

The first result I want to show you is striking. Panel A of **Figure 1** indicates that the CFOs' one-year *ex ante* risk premiums (framed in the survey as the excess return of stocks over U.S. T-bills) vary considerably over time. The last survey, finished on September 10, 2001, indicates the CFOs were forecasting at that time a one-year-ahead risk premium of, effectively, zero. The 10-year-horizon *ex ante* risk premium, given in Panel B, is interesting because it is higher than the 1-year-horizon forecast and is stable from survey to survey at about 4 percent (400 bps). Note that the September 10, 2001, forecast is 3.6 percent.

Figure 1. Survey Respondents' One-Year and Ten-Year Risk Premium Expectations

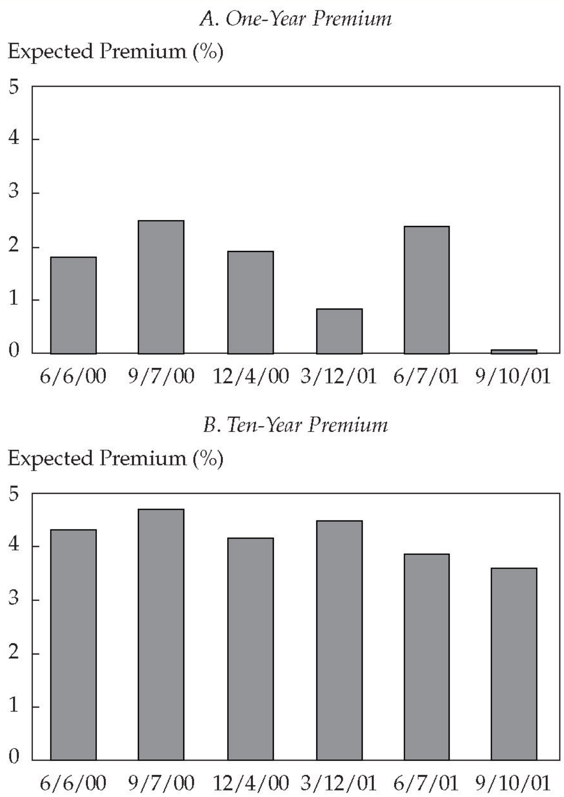

One of the first aspects we investigate is whether the CFOs' expectations about future returns are influenced by past returns. That is, if the market has performed poorly in the immediate past, does this performance lead to lower expected returns? **Figure 2** is a simple plot of the expected one-year equity risk premium against the previous quarter's return. (As we go through the analysis, please keep in mind that one can really be fooled by having so few observations. Indeed, this problem is exactly the reason we chose to present most of the results graphically. By eyeballing the data, you can see whether one observation is driving the relationship.) Figure 2 shows a fairly reliable positive relationship between past return and future near-term expected risk premium. Also, we found that you can pull out any of these observations and the fit is still similar. Apparently, a one-year-horizon forecast carries what Graham and I call "expectational momentum." Therefore, negative returns influence respondents to lower their forecast of the short-term future premium.

Figure 2. One-Year Risk Premium and Recent Returns

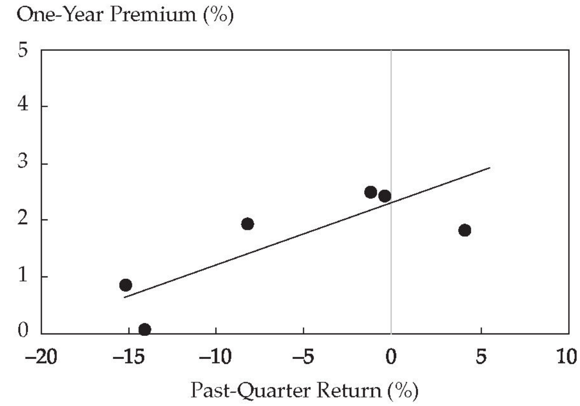

Notes: $y = 0.1096x + 2.3068$; $R^2 = 0.7141$.

Figure 3 plots the same variables for the 10-year horizon. There is a slight positive relationship between the past quarter's return and the *ex ante* 10-year-horizon risk premium, but it is not nearly as positive as the relationship observed for the 1-year horizon.

Figure 3. Ten-Year Risk Premium and Recent Returns

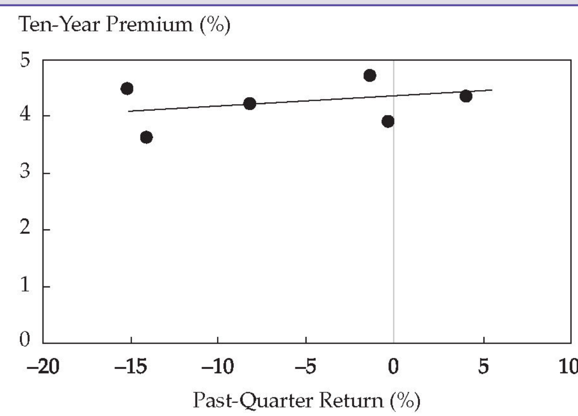

Notes: $y = 0.0179x + 4.3469$; $R^2 = 0.1529$.

We measured expected market volatility by deducing each respondent's probability distribution. We asked the respondents to provide a high and a low forecast by finishing two sentences: "During the next year, there is a 1-in-10 chance the S&P 500 return will be *higher* than ______ percent" and "During the next year, there is a 1-in-10 chance the S&P 500 return will be *lower* than _______ percent." The expected market volatility is a combination of the average of the individual expected volatilities (which I will refer to in the figures as "average volatility") plus the dispersion of the risk premium forecasts (referred to as "disagreement").[2]

Figure 4 shows that (annualized) average expected volatility for the one-year horizon is weakly negatively related to the past quarter's return. In fact, if one observation were pulled out, we might find no relationship whatsoever. And **Figure 5** shows the (annualized) disagreement component—basically, the standard deviation of the risk premium forecast—for the one-year horizon. The disagreement component for the one-year horizon is strongly related to the past quarter's return. A bad past return suggests a higher disagreement volatility. Even with so few data points, this relationship appears to be strong.

Figure 4. Average (One-Year-Horizon) Volatility and Recent Returns

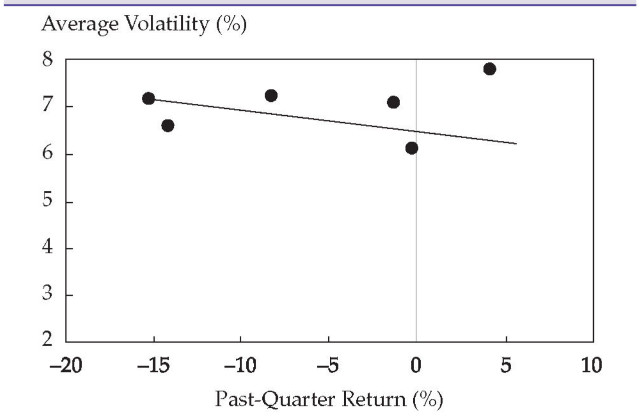

Notes: $y = -0.0452x + 6.4722$; $R^2 = 0.1282$.

One thing to keep in mind is that these points on Figures 4 and 5 are annualized. When you examine the individual volatilities, you find that these respondents are extremely confident in their assessments. The result is a 6–7 percent annualized volatility in

Figure 5. Disagreement (One-Year Horizon) Volatility and Recent Returns

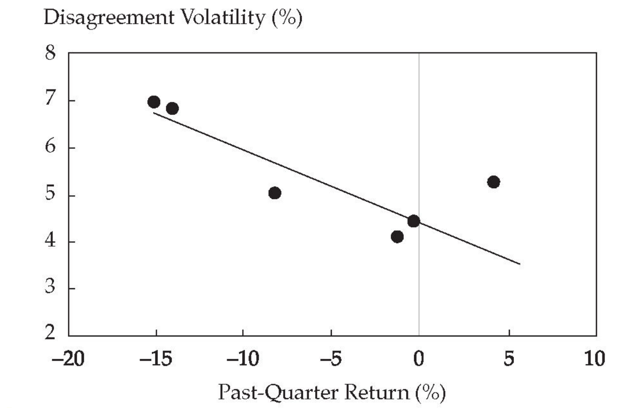

Notes: $y = -0.153x + 4.3658$; $R^2 = 0.7298$.

the one-year-horizon *ex ante* risk premium. This volatility is much smaller than typical market estimates, such as the Chicago Board Options Exchange VIX (Volatility Index) number on the S&P 100 option, which averages around 20 percent.

We also found that our measure of asymmetry is positively related to the past quarter's return. Given that we get the tails of the distribution, we can look at the mass above and below the mean and compare them, which gives us an *ex ante* measure of skewness. If past returns are negative, we find more negative *ex ante* skewness in the data.

Instead of looking at the relationship of the forecasted risk premium to past return, **Figure 6** relates the forecasted (*ex ante*) risk premium to expected (*ex ante*) volatility. Many papers in academic finance have examined the relationship between expected risk and expected reward. Intuitively, one would expect the

Figure 6. Expected Average Volatility and Expected Risk Premium: One-Year Horizon

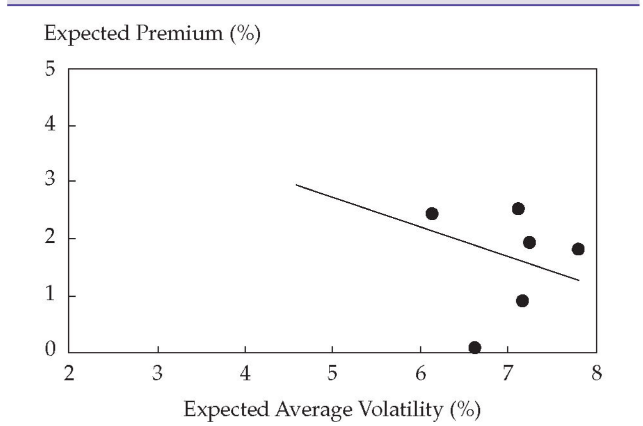

Notes: $y = -0.5178x + 5.2945$; $R^2 = 0.2538$.

[2] Market volatility was measured as

$$\text{var}[r] = E[\text{var}(r|Z)] + \text{var}[E(r|Z)],$$

where r is the market return, Z is the information that the CFOs are using to form their forecasts, $[E(r|Z)]$ is the expected risk premium conditional on the CFO's information, $E[\text{var}(r|Z)]$ is the average of each CFO's individual volatility estimate, and $\text{var}[E(r|Z)]$ is disagreement volatility or the variance of the CFOs' forecasts of the premium. Individual volatilities were measured as

$$\text{var} = \left[\frac{x(0.90) - x(0.10)}{2.65}\right]^2,$$

where $x(0.90)$ is the "one in ten chance that the return will be higher than" and $x(0.10)$ is the "one in ten chance that the return will be lower than." The equation for individual volatilities is from Davidson and Cooper (1976).

relationship to be positive, but the literature is actually split. Indeed, many papers have documented a negative relationship, which is basically what we see for the one-year-horizon predictions. In Figure 6, the *ex ante* premium and the *ex ante* average volatility appear to be weakly negatively related. **Figure 7** plots the one-year-horizon expected risk premium against disagreement about the expected premium. The result is a strongly negative relationship: The higher the disagreement, the lower the expected premium over one year. Again, almost any observation could be pulled out without changing the degree of fit.

Figure 7. Disagreement Volatility and Expected Risk Premium: One-Year Horizon

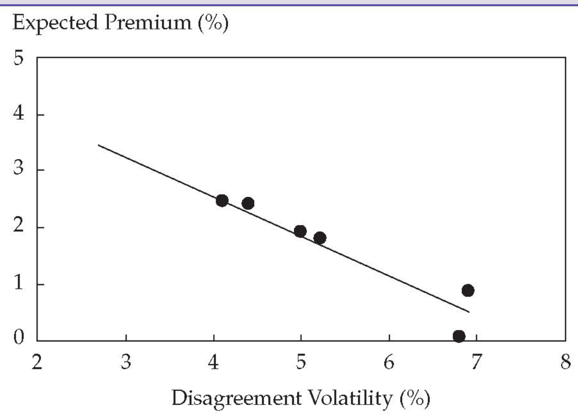

Notes: $y = -0.6977x + 5.3410$; $R^2 = 0.9283$.

Using the same variables as in Figure 7 and keeping the scale the same, **Figure 8** shows the data for the 10-year horizon. The fit is again strikingly good, but the relationship is positive. Notice that the disagreement is much smaller for the 10-year horizon than for the 1-year horizon. This positive relationship between the *ex ante* premium and *ex ante* volatility is suggested by basic asset-pricing theory.

The latest survey documented in Figures 2–8 is June 1, 2001, plus data returned to us by September 10, 2001. We just happened to fax our most recent quarterly survey to the survey participants at 8:00 a.m. on the morning of September 10. I did not include observations from the surveys returned on September 11 because the survey might have been completed on either September 10 or 11, and classification of the responses as pre- or post-September 11 was not possible. The response data we received on September 12 or later we maintained and analyzed separately. **Table 1** provides a comparison of pre- and post-September 11 data for the 1- and 10-year horizons. Although the size of the sample is small (33 observations), one can see the impact of September

Figure 8. Disagreement Volatility and Expected Risk Premium: Ten-Year Horizon

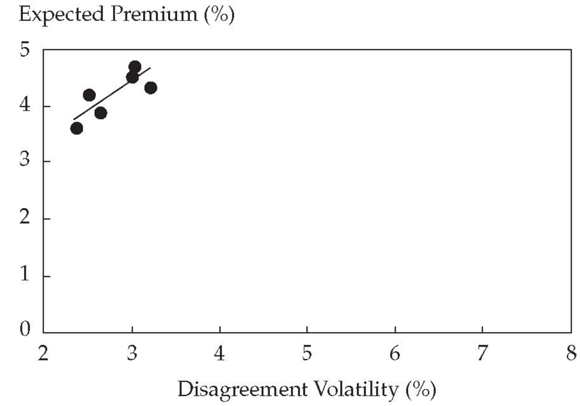

Notes: $y = 0.9949x + 1.4616$; $R^2 = 0.6679$.

11. The 1-year-horizon mean forecasted premium decreases after September 11, but volatility—both disagreement and average—increases. For the 10-year horizon, the mean forecasted premium and disagreement volatility increase. I'll be the first to admit that these results are not statistically significant, but the data tell an interesting story. After September 11, perceived risk increases—which is no surprise. In the short term, participants believe that market returns will be lower. In the long term, however, premiums increase to compensate for this additional risk.

Table 1. Impact of September 11, 2001: Equity Risk Premium and Volatility

Measure	Before	After
Observations	127	33
1-year premium		
Mean premium	0.05 %	–0.70 %
Average volatility	6.79	9.76
Disagreement volatility	6.61	7.86
10-year premium		
Mean premium	3.63 %	4.82 %
Disagreement volatility	2.36	3.03

Implications of Results

So, what have we learned from this exercise? First, expectations are affected, at least in the short term, by what has happened in the recent past—an expectational momentum effect. Second, these new expectational data appear to validate the so-called leverage effect—that negative returns increase expected volatility. Third, the individual volatilities (at 6–7 percent) seem very low, given what we would have expected. And fourth, there is apparently a positive relationship between risk and expected return (or the risk premium) only at longer horizons. So, the horizon is critical.

How should we interpret these results, what are the outstanding issues, and where do we go from here? The CFOs in the survey are probably not using their one-year expected risk premiums for one-year project evaluations. What CFOs think is going to happen in the market is different from what they use as the hurdle rate for an investment. I do think that the 10-year-horizon risk premium estimates we are getting from them are close to what they are using. An interesting paper being circulating by Ravi Jagannathan and Iwan Meier (2001) makes some of these same arguments—that higher hurdle rates are probably being used for a number of reasons: the scarcity of management time, the desire to wait for the best projects, and financial flexibility. Corporate managers want to wait for the best project, and with limited management time, a hurdle rate that is higher than what would be implied by a simple asset-pricing model allows that time.

Another angle is that the premium *should* be high in times of recession. Indeed, a lot of research documents apparently countercyclical behavior in the premium. Such behavior implies that today's one-year-horizon investment should have a high hurdle rate.

Further Research

We hope our research sheds some light on the measure of expectations. I believe in asset-pricing models based on fundamentals, but it is also enlightening to observe a direct measure of expectations. Our data may not be the true expectations, but they supply additional information about the *ex ante* risk premium in terms of investment horizon, expected volatility, and asymmetry.

Our next step is to conduct interviews in the first week of December 2001 with a number of the CFOs participating in the multiperiod survey. We have already carried out a few preliminary interviews, and we find it extraordinary how much thought CFOs have given to these issues. The main question we want to ask in December is the reason (or reasons) for the difference between their risk premium forecasts for a one-year horizon and the actual internal hurdle rates they use to evaluate one-year-horizon projects. How do CFOs use the *ex ante* risk premium in terms of making real allocation decisions? I will keep you updated on the progress of our research project.

Implications for Asset Allocation, Portfolio Management, and Future Research II

Campbell Harvey
Duke University, Durham, North Carolina
National Bureau of Economic Research, Cambridge, Massachusetts

SUMMARY

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

The presentation made by Campbell Harvey was unique, in that it was based essentially on surveys of investor expected risk premiums. What he had heard from the previous speakers was how difficult it is to get a measure of investor expectations.

Harvey's surveys, over time, of chief financial officers offered what he considered to be a less biased sample than the surveys that have been made of economists or financial analysts. CFOs are known to be concerned about a measure of their cost of capital for investment planning purposes and have no reason to favor high or low forecasts. He stated that, although he does not see the survey results as a replacement for the kind of analyses presented by previous speakers, he does believe that the surveys add valuable information.

The survey questions and responses were for 1-year and 10-year time horizons, which provided an opportunity to compare short-term with long-term expectations. The surveys elicited information not only on the expected premiums but also on the probability distributions of the respondents' forecasts. Harvey considered two components of expected market volatility: the average of the individual expected volatilities (from each individual's probability distribution) and the disagreement over the risk premium forecasts (the standard deviation of the risk premium forecasts).

Figure 1 shows the results of six surveys asking for a 1-year risk premium estimate and a 10-year estimate. The 10-year forecasts show little variation, whereas the 1-year forecasts vary widely through time. The 10-year forecasts are also consistently higher than the 1-year forecasts.

Figure 1. Survey Respondents' One-Year and Ten-Year Risk Premium Expectations

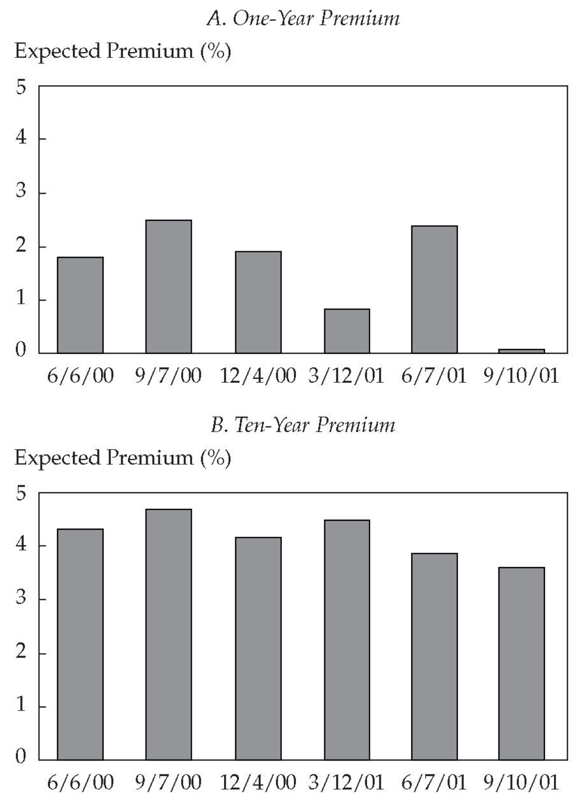

Figure 2 shows the influence of past returns on forecasts of 1-year premiums, and **Figure 3** does the same for 10-year premiums. Past returns had a positive impact on 1-year forecasts and a very slight positive effect on 10-year forecasts. Past returns also had a weak negative effect on expected 1-year average volatility and a strong negative effect on disagreement. They had a strong positive effect on expected skewness. Negative returns led to more negative skewness in the forecasts.

Turning to the effect of expected rather than past returns, Harvey showed in **Figure 6** that the average

Figure 2. One-Year Risk Premium and Recent Returns

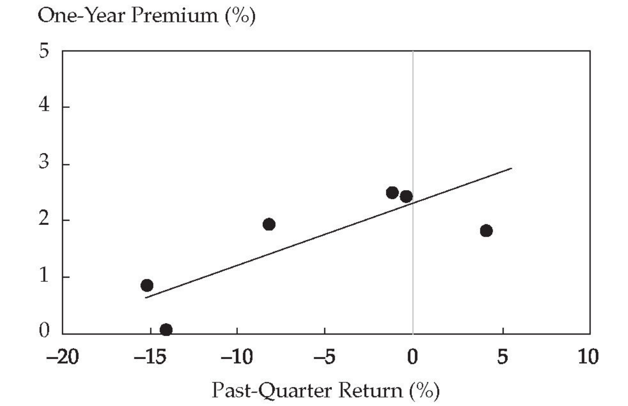

Notes: $y = 0.1096x + 2.3068$; $R^2 = 0.7141$.

Figure 3. Ten-Year Risk Premium and Recent Returns

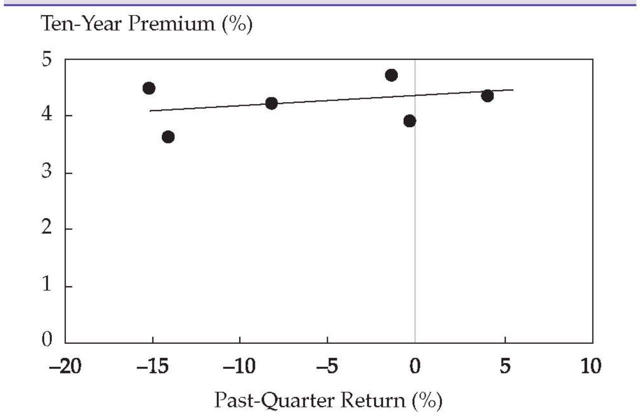

Notes: $y = 0.0179x + 4.3469$; $R^2 = 0.1529$.

of individual volatilities is weakly *negatively* related to expected 1-year returns.[1] One-year expected returns were found to be strongly negatively related to disagreement volatility, as shown in **Figure 7**. This finding may seem counter to the usual risk–expected return theories, but the finding is for very short term forecasts. For the 10-year horizon shown in **Figure 8**, however, expected returns are strongly *positively* related to disagreement—which is consistent with the way we usually think about risk and expected reward.

Harvey reported the impact of the events of September 11, 2001, in **Table 1**. After the crisis, the CFOs revised expected returns for the 1-year forecasts downward. For both the 1-year and the 10-year forecasts, expected volatility increased after the crisis.

[1] Table and figure numbers in each Summary correspond to the table and figure numbers in the full presentation.

Figure 6. Expected Average Volatility and Expected Risk Premium: One-Year Horizon

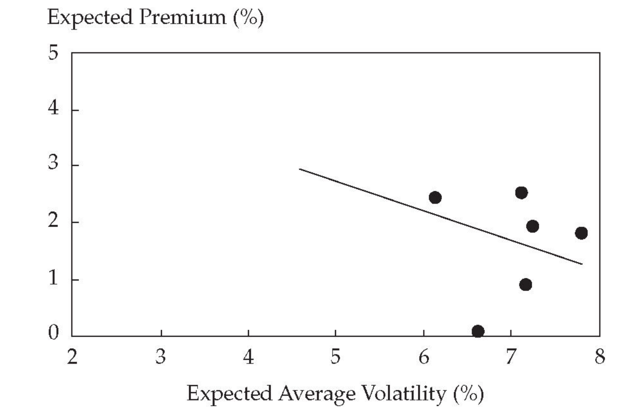

Notes: $y = -0.5178x + 5.2945$; $R^2 = 0.2538$.

Figure 7. Disagreement Volatility and Expected Risk Premium: One-Year Horizon

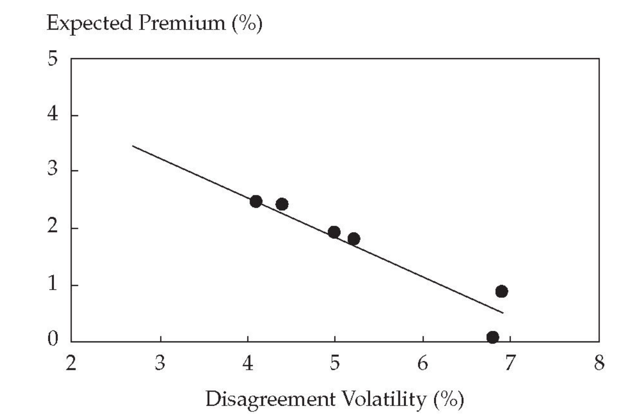

Notes: $y = -0.6977x + 5.3410$; $R^2 = 0.9283$.

Figure 8. Disagreement Volatility and Expected Risk Premium: Ten-Year Horizon

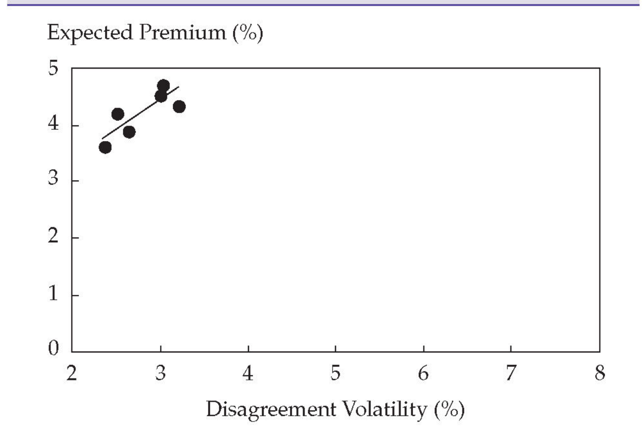

Notes: $y = 0.9949x + 1.4616$; $R^2 = 0.6679$.

Table 1. Impact of September 11, 2001: Equity Risk Premium and Volatility

Measure	Before	After
Observations	127	33
1-year premium		
Mean premium	0.05%	–0.70%
Average volatility	6.79	9.76
Disagreement volatility	6.61	7.86
10-year premium		
Mean premium	3.63%	4.82%
Disagreement volatility	2.36	3.03

Summarizing, Harvey presented the following conclusions:

- Survey measures of expectations provide useful alternatives to statistical measurements.
- Return forecasts are positively influenced by past returns—what John Graham and Harvey (2001a) call "expectational momentum."
- Expected volatility is negatively related to past returns.
- Individual volatilities seem very low; the respondents seem very confident in their forecasts.
- Time horizon makes a big difference. There is a positive relationship between risk and expected return but only for long-horizon forecasts.

In closing, Harvey expressed doubt that the CFOs were actually using their 1-year forecasts for hurdle rates in 1-year project evaluations. He suggested that there is a difference between what CFOs believe will happen to the market next year and the rate of return they would accept for a new project. The 10-year forecasts are probably closer to what the CFOs are using for the cost of capital.

Implications for Asset Allocation, Portfolio Management, and Future Research: Discussion

Roger Ibbotson (*Moderator*)
Robert Arnott
John Campbell
Bradford Cornell
William Goetzmann
Campbell Harvey
Martin Leibowitz
Thomas Philips
William Reichenstein, CFA

ROGER IBBOTSON (Moderator)

I was particularly pleased to see Campbell Harvey's paper because we have seen surveys of financial analysts, individuals, and economists (such as Welch's 2000 survey of financial economists), but the Graham and Harvey (2001a, 2001b) survey breaks new ground by surveying a particularly astute group. The results of their survey bring fresh information to the table. The survey was also well designed, which gives us confidence in the data.

I think each of us understands that we are concerned with equity risk premiums looking forward, but the distance we are looking ahead, our horizons, may differ. And today we have had both discussions—looking short term and looking out long term. The differences between the short-run and the long-run risk premium were certainly brought out by Rajnish Mehra [in the "Current Estimates and Prospects for Change" session] and are highlighted in the Graham and Harvey work.

I would like to present a few ideas from a paper that Peng Chen and I wrote (Ibbotson and Chen 2002) that uses much of the same data that Rob Arnott used but interprets the data almost completely differently. One of the reasons for the lack of overlap in interpretations is that Rob's primary focus is a short-run prediction of the market.

Figure 1 is yet another P/E chart—this one based on the Wilson and Jones (forthcoming 2002) data because their earnings data match the S&P 500 Index earnings data. The S&P 500 had very low, not negative but very low, earnings in the 1930s, and the actual maximum P/E is off the chart for that period. Figure 1 begins with a P/E, calculated as price divided by prior-year earnings, of 10.22 in 1926 and ends with a P/E of 25.96 at year-end 2000 (the October 2001 P/E, excluding extraordinary earnings, is 21); that growth from about 10 to the most recent P/E is an important consideration in the forecast I will discuss.

The forecast that Peng and I are making is based on the real drivers of P/E growth. We focus on the contribution of earnings to P/E growth and on GDP. **Table 1** shows the historical average nominal return for stocks over the 75-year period of 1926 through 2000 to be 10.70 percent. We can break that nominal stock return into its contributing components: about 3 percentage points (pps) inflation, and so forth. The P/E growth rate from a multiple of about 10 in 1926 to a multiple of almost 26 in 2000 amounts to 1.25 percent a year. When we make our forecasts, we remove that historical growth rate because that P/E jump from 10 to 26, in our opinion, will not be repeated. The "Earnings Forecast" column in Table 1 shows what history was without the P/E growth rate; that is, the forecasted return is 1.25 pps less than the historical return.

Table 1. Historical and Forecasted Components of Stock Returns, 1926–2000

Component	Historical[a]	Earnings Forecast
Income	4.28 pps	4.28 pps
P/E growth	1.25	—
Earnings growth	1.75	1.75
Inflation	3.08	3.08

[a]Total historical return for the period is 10.70 percent; data do not sum to that total because of the geometrical mathematics used.

Figure 2 provides the historical growth of per capita GDP and of earnings, dividends, and capital gains on a per share, not aggregated, basis. All are indexed to $1 at the end of 1925. The capital gains grow to about $90 at the end of 2000—the most growth of any of the measures shown. Earnings are less because of the increase in the P/E multiple. The $90 is the $36 multiplied by 2.5, which was the P/E

Figure 1. The P/E, December 1925–December 2000

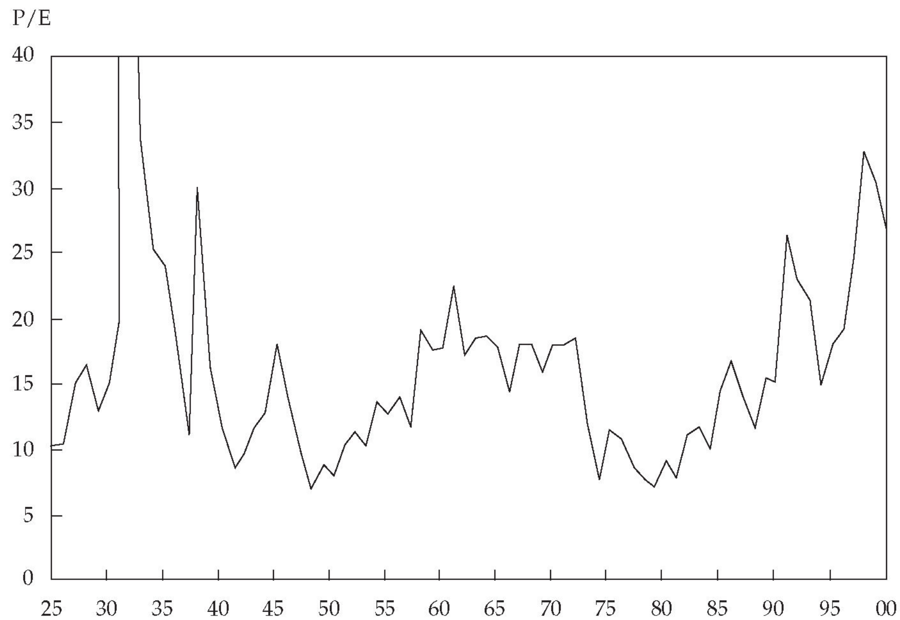

Note: The P/E for December 1932 was 136.5.

Figure 2. Historical Growth of per Capita GDP and of per Share Earnings, Dividends, and Capital Gains, December 1925–December 2000

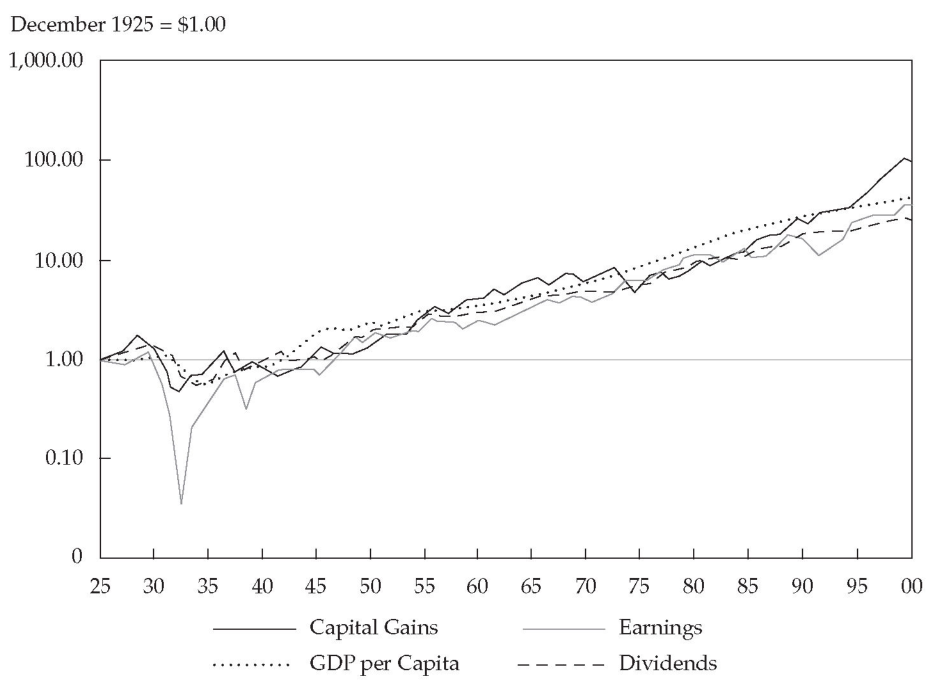

Note: At end date, capital gains were $90.50, GDP per capita was $44.10, earnings were $35.60, and dividends were $24.20.

change from 10 to 26. The line for GDP per capita shows that the economy (on a per capita basis) has outgrown earnings by a small amount over the entire period. And finally, the growth in dividends trails the pack. So, I very much agree with the comment that Bill Reichenstein made earlier today that dividends are not a good forecasting tool; they grow the most slowly and even distort the picture for earnings growth [see "Current Estimates and Prospects for Change: Discussion"].

I am struck by how tied together each data series is—how the stock market is related to the economy, which is related to earnings, which are related to dividends. Although the link between earnings and dividends is a little less close than the other links, it is still there. One of the reasons Peng and I wanted to carry out this type of analysis is that the economy *should* be reflected in the stock market. And in fact, the separation in their behaviors is solely the result of the changing P/E, which we have thus removed from our forecasts. The P/E rose from 1926 to 2000 for a reason, but that reason will not continually recur in perpetuity. For that annual growth rate in the P/E multiple of 1.25 percent a year to continue, to assume that it will replicate, would mean that in another 75 years, the P/E will have grown to 62.

Figure 3 shows why dividends are not a good tool for forecasting the future. Dividend yields started the period at 5.15 percent and averaged 4.28 percent over the past 75 years; if you include the data for the 19th century, the historical average dividend yield is much higher. Every time we found a dividend for the 19th century, it seemed to be 100 percent. The dividend yield has now dropped to 1.10 percent (the most recent year would push it up somewhat). Thus, a long-run secular decline has occurred in the dividend yield, which was largely caused by the decreasing payout ratio. As **Figure 4** shows, the payout ratio, which began the period at 46.68 percent and averaged almost 60 percent over the 1926–2000 period, is now 31.78 percent.

Several reasons could explain the trend toward lower payout ratios. We interpret the trend as an issue of trust and changing attitudes about trust. As investors place more trust in the companies in which they invest and in the financial market system, shareholders no longer require that the companies pay all of their earnings to the shareholders; the discipline that dividends were designed to impose on corporations is gradually falling by the wayside. Another possible reason for the trend toward lower payout ratios is that, of course, dividends and capital gains (the fruit of reinvested corporate earnings) are taxed differently—providing an incentive for shareholders to relax their desire for company earnings to be paid out as dividends. Moreover, today, earnings can be taken out in many forms, such as share repurchases, buyouts in a merger or acquisition, or investment in internal projects of a company. I predict that these myriad forms of paying out earnings will remain. A

Figure 3. Dividend Yield, December 1925–December 2000

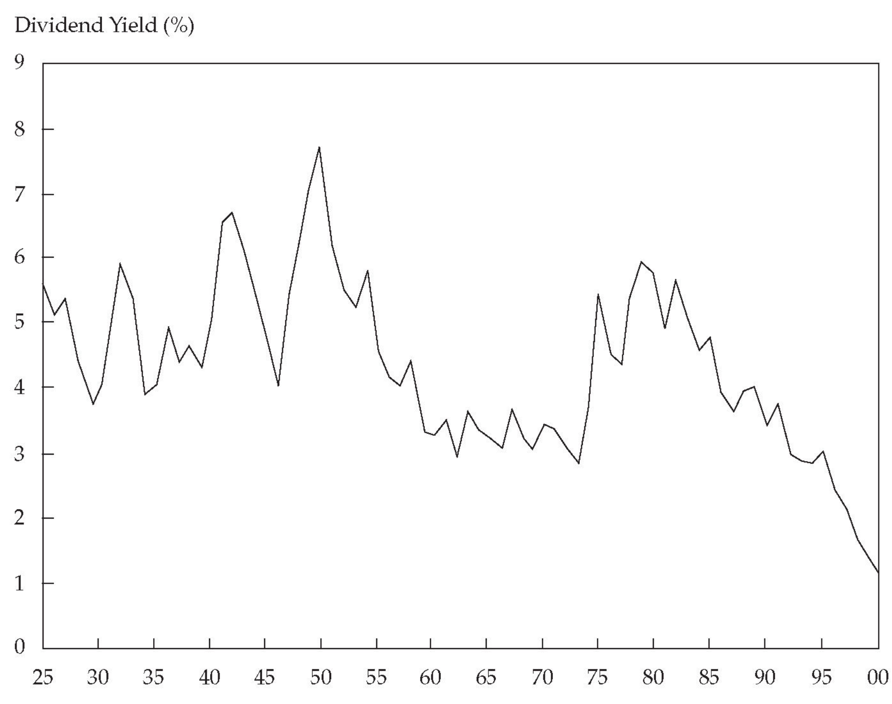

Figure 4. Dividend Payout Ratio, December 1925–December 2000

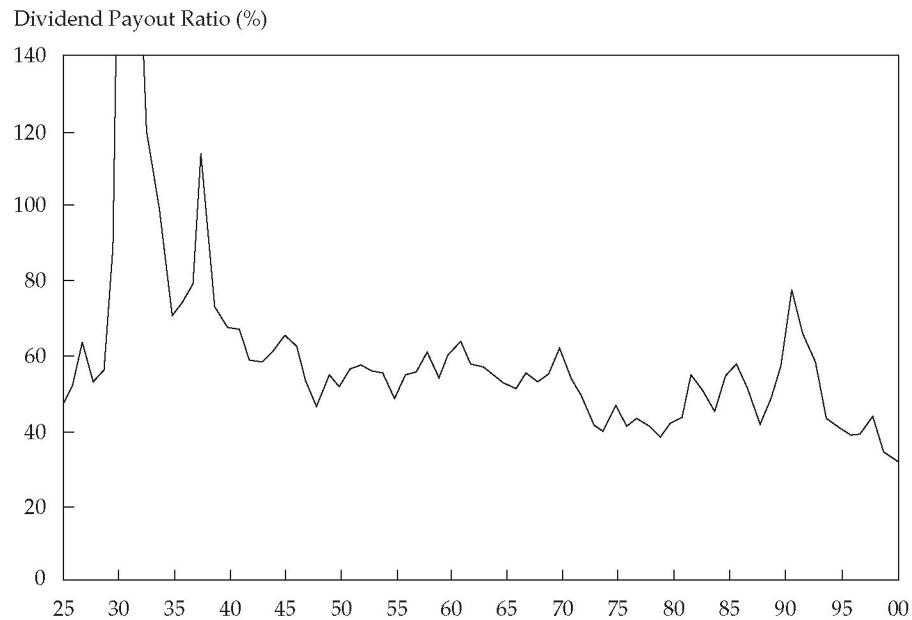

Note: The payout ratio as of December 1931 was 190.52 percent; as of December 1932, it was 929.12 percent.

larger and larger portion of companies in the market are not paying earnings out in the form of dividends. For example, the technology companies do not pay out any of their earnings as dividends. Thus, the payout ratio is not stable, and we may see it continue to fall.

A contender in the race to be a reliable forecasting tool (one that a number of people have already discussed today) is the dividend yield model in one of its many forms. If you could accept the dividend yield model by itself and with its purest assumptions—that is, the dividend yield plus dividend growth, assuming constant growth—the model would be a forecast of the stock market. But there are three problems with the pure dividend yield model that we must make adjustments for if the model is to be useful for forecasting. The first two problems are potential violations of Modigliani and Miller theory.

I am assuming that M&M holds true. (Despite what some of you have said about how dividend payouts do not seem to be reinvested in anything at all, I am clearly on the other side of that argument. If there is any truth to that supposition, however, that theory needs further investigation.) So, the first problem with some forms of the dividend yield model is that they violate M&M because they assume you can add the current dividend yield (which is now 1.10 percent) to historical dividend growth. Historical dividend growth underestimates historical earnings growth, however, because of the decrease in the payout ratio. Dividends have run slowest in the growth race because the payout ratio has continually dropped.

The second problem with using the dividend yield model as a forecasting tool (and it is, again, a violation of M&M) is that if the low payout ratios of today (31.8 percent) were reflected in the historical series, the percentage of earnings retained would have been higher and, therefore, historical earnings would have grown faster than observed. In short, the first problem is that dividend growth has been too slow historically, and the second problem is that with further earnings retention, historical earnings growth would have been potentially faster than observed.

The third problem with the dividend yield approach is the high P/E multiple observed today—over 25. Unlike some of you, I am going to assume efficient markets, which in this case I take to mean that the current high P/E implies higher-than-average future EPS growth.

My estimate of the average geometric equity risk premium is about 4 percent relative to the long-term bond yield. It is, however, 1.25 percent lower than the pure sample geometric mean from the risk premium of the Ibbotson and Sinquefield study (Ibbotson Associates 2001).

We have had some debate today on future growth rates—specifically for the 10-year horizon. Data that Peng and I are studying provide some support for the tie between high P/Es and high future growth. One of the problems with the 10-year horizon is that 10 years is not really long enough to encompass many independent events.

The extreme end of the spectrum of proponents of the dividend yield model would support using past dividend growth to forecast future dividend growth, then add current income. (Of course, that method almost wipes out the risk premium, and in some ways, it is actually similar to what Rob Arnott presented.)

In our response, we make three adjustments to the dividend yield model shown in the third column ("Current Dividend Forecast") of **Figure 5**. These are shown in the fourth column ("Current Dividend Forecast with Additional Growth"). We add 0.51 pp so that historical dividend growth matches historical earnings growth, we add an additional 0.95 pp because of the extra retention associated with the current record low payout rate, and finally we add 2.28 pps to future earnings growth to reflect the current high P/E that we assume forecasts higher earnings growth.

What about long-term earnings growth? Corporate America is likely to proceed in the next quarter century as it did in the previous 75 years. Corporate cash will be used for projects, investments, share repurchases, and acquisitions, but less and less will it be used for dividend payouts. Future earnings growth will be higher than past growth because of lower dividend payouts and the high current P/E. For the next 25 years, I predict (1) stocks will outperform bonds, (2) increased earnings growth will offset future low dividend yields, (3) the P/E jump from 10 to 26 will not repeat, and (4) the stock market return will provide more than 9 percent a year over the 25-year period.

JOHN CAMPBELL: When you make the adjustments, aren't you assuming not only efficient markets but also a constant discount rate? If so, you are assuming the answer. We are trying to find out what the discount rate is, but you assume the discount rate in your calculation. If so, aren't you bound to come up with an answer for the end that is the same as historical norms going in?

Figure 5. Historical versus Forecasts Based on Earnings and Dividend Models

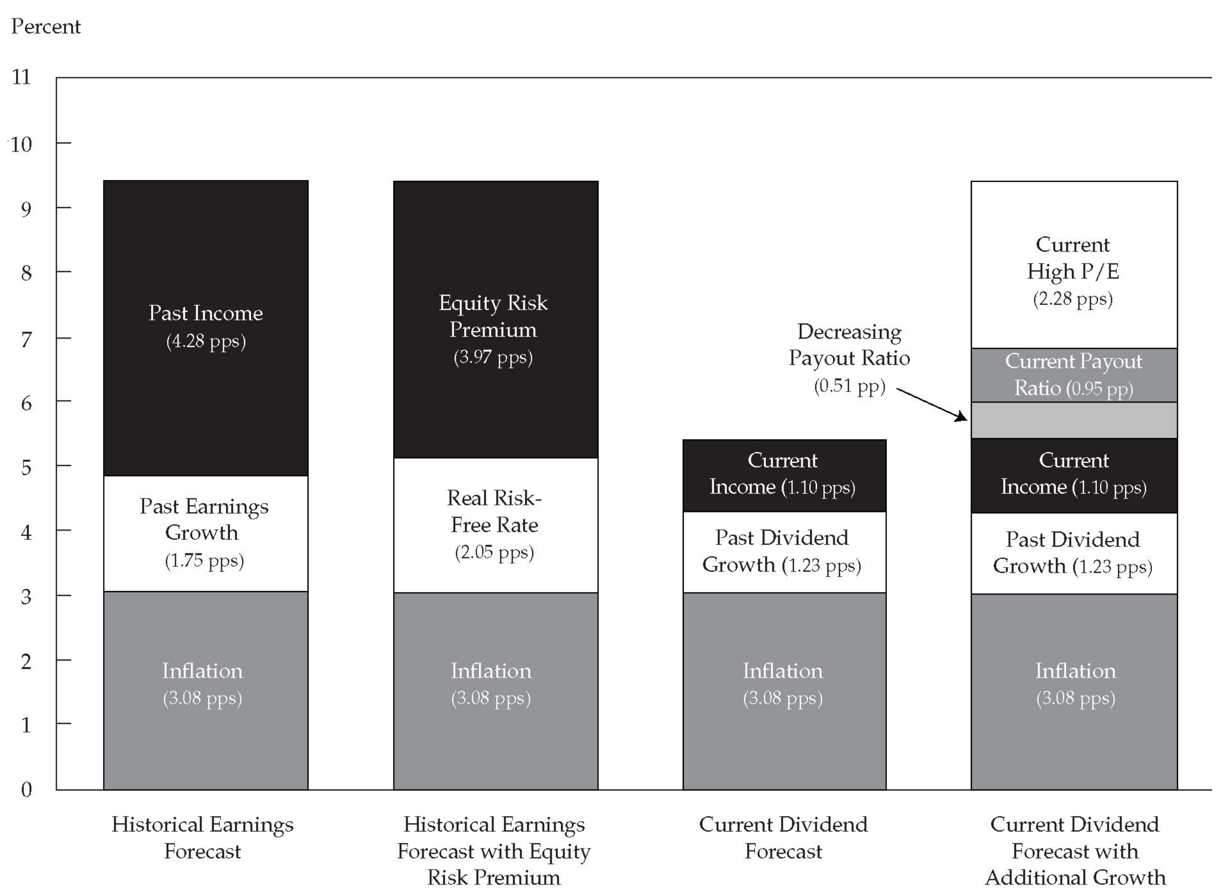

IBBOTSON: True. In addition to assuming an efficient market (M&M), we are not assuming that the discount rate is dynamic. We are assuming it to be unknown, and we are searching for the single discount rate that best describes history. The presumption is that history can be extrapolated forward. It could be considered a reconciliation between the two approaches. Certainly, our quest is debatable.

BRADFORD CORNELL: I have some questions for Campbell Harvey. Are CFOs really not using their one-year-horizon market forecasts in evaluating their internal investments? Maybe the one-year market forecast they provide you is just a throw-away number; they are so uncertain about it that they do not incorporate it into any decision they make. If they really believe that the equity risk premium is zero today, shouldn't they be issuing stock?

CAMPBELL HARVEY: I think this survey gives us respondents' guesses of what is going to happen in the market; it does not necessarily map into what they are going to do in terms of their real project evaluations at a one-year horizon. In a recent working paper by Jagannathan and Meier (2001), which is based on some older work by McDonald and Siegal (1986), they say people tend to have higher hurdle rates than what the capital asset pricing model (CAPM) would suggest. CFOs are looking for the best projects, internal investments that throw off the best return, and there is no way they are going to accept a project with a rate of return equal to the T-bill rate—even if they expect next year's market return to be basically the same as the T-bill's return. So, what the data suggest to me is that there is a big difference between the short-horizon expectation of return and the hurdle rate one would actually use in terms of project evaluation. Of course, I want to go deeper into this problem by asking the survey participants for more details.

ROBERT ARNOTT: One would assume that to arrive at the estimated required return of any new commitment, a "credibility" hurdle rate is added on top of the cost-of-capital hurdle rate. Those cost-of-capital hurdle rates are always optimistic, so the credibility rate is added and is part of where the reported hurdle rate in the responses comes from.

MARTIN LEIBOWITZ: Just one clarification: How did your 10-year risk premium, 4.5 percent, relate to the hurdle rate? Do you have any evidence of what that longer-term hurdle rate is?

HARVEY: For the 10-year horizon, the risk premium reported is closer to the hurdle rate for internal projects than for the 1-year horizon. We don't have much information about the longer-term hurdle rate, but the next phase of my research with John Graham will be interviewing the CFO participants to shed additional light on these issues.

WILLIAM GOETZMANN: I was very excited to see Campbell Harvey's paper—to see more interesting data about dispersion of opinion. I know that in one of your earlier papers—the one on the market-timing ability of investment newsletter writers (Graham and Harvey 1996)—you unexpectedly found dispersion of opinion that had some forecasting ability. Cragg and Malkiel (1982) also found some dispersion in analysts' forecasts in relation to risk. Also, Massimo Massa and I have been finding some information about dispersion related to price effects and so forth (Goetzmann and Massa 2001). What particularly strikes me in looking at your results is the consistent message that this dispersion of opinion is having interesting effects that we ought to explore. If you are going to be talking to these CFOs, it would be great to find out more about the basis for the dispersion. It is an interesting potential area of research.

HARVEY: We have a lot of data on earnings forecasts, but I am more interested in the dispersion than the actual forecasts. An older paper by Frankel and Froot (1990) looked at dispersion of beliefs in terms of currency forecasting. It is very impressive. So, I agree that this area is worthy of more research.

THOMAS PHILIPS: I want to address the question about forecasts versus hurdle rates by describing an experience that I had. When I talk to our corporate clients, I often ask if they need help estimating their cost of capital (which, of course, is the same as the expected return) and I ask how they do it currently. Some tell me that they use the CAPM, while others say they use a more complicated factor model. But one answer stands out for its simplicity and its brilliance. At National Service Industries, an executive told me that his cost of capital was 10 percent. I asked him how he knew that it was 10 percent. He replied that he did not *know* that it was 10 percent. So, I queried further: "Why, then, do you assert that it is 10 percent?" He replied, "In my world, the cost of capital is not very important in terms of making new investment decisions. We have a hurdle rate to make that type of decision. The cost of capital is important to us because the lines of business that we are in are not fabulously profitable, and the simplest mistake we can make is to squander the capital we have invested in them. The one thing I want to do is to have every employee understand that capital is a real input and that it is incredibly easy to squander. When I use 10 percent as the cost of capital, everyone from the janitor to the CEO can apply it. They can move a decimal point; they can divide by 10. So, I can explain to them in simple terms that $1 million worth of equipment sitting idle represents $100,000 of real money going down the tubes every year. And that ability is much more important to me and to the company than having the right answer." Theoretically, he has the wrong answer, but in spite of that, his answer and approach are absolutely brilliant.

The other comment that I want to make is an observation on the difference in earnings growth rates. Roger Ibbotson is showing it growing close to per capita GDP.

ARNOTT: No, he has it growing faster than GDP.

PHILIPS: Roughly the same rate.

IBBOTSON: Historically, it is the same.

ARNOTT: But now the payout ratio is lower, so earnings would have to grow faster. Earnings growth is going to gain on GDP on a per share basis, not necessarily on an aggregate basis as Bradford Cornell was talking about.

WILLIAM REICHENSTEIN: Going back to what Rob Arnott said about taking another look at tactical asset allocation. Let's say that over the next 10 years, stocks, bonds, and cash will all produce a 10 percent rate of return. It seems to me the 10-year return should not make any difference; the asset-allocation decision is relatively insignificant at that point.

ARNOTT: Correct, the policy asset allocation decision is insignificant. For rebalancing to add value, for tactical asset allocation to add value, the absolutely crucial premise is that reversion to the mean will occur in at least a weak form.

REICHENSTEIN: That is when you pick up your alpha?

ARNOTT: Right. The presumption is based on a long-term historical record for live TAA experience. Even when it did not add value (in the 1990s), it did produce alpha. If there were not some weak reversion to the mean at work in the 1990s, it would not have produced an alpha.

LEIBOWITZ: Why do you say policy allocation is invariant? Even if you have zero difference in returns, you still have volatility.

ARNOTT: I am assuming geometric, not arithmetic, returns. If we assume arithmetic returns are the same, then the volatility differences carry a cost. If we assume the geometric returns are the same, then the return-maximizing portfolio is the risk-minimizing portfolio, which would probably have an allocation of only 10–20 percent equities. But the difference in returns would be tiny, so whether the allocation was 20/80 or 80/20 would not make much difference in the return.

LEIBOWITZ: But you would not have much in equities?

ARNOTT: This message is not welcomed with open arms by investors or investment practitioners. It has not been good for First Quadrant's business for me to publish this sort of stuff. Some consultants are annoyed because we are saying, basically, that the assumptions they are endorsing are wrong. Clients don't want to hear it because we've been correct for the last year and a half, and the losses hurt. When we first proposed the idea, it was viewed as slightly flaky, but since then, it's been on target—which has made some people even angrier.

GOETZMANN: I'm a bit confused. Are you talking about just *your* track record or evidence about TAA in general? I haven't seen any empirical evidence indicating that, on average (or even in the tails), any tactical allocators have been successful.

ARNOTT: I am speaking on the basis of our track record and what little information I can garner about competitors' track records. The comparative studies, like the one that Tom Philips did (Philips, Rogers, and Capaldi 1996), have dwindled to next to nothing because no one is interested in TAA. Our founding chairman was fond of saying, "Don't buy what's easy to sell. Do buy what's tough to sell." Well, TAA is tough to sell right now. I think it is an interesting idea that has fallen from favor in a circumstance where, prospectively, it is probably going to produce the kind of results that we had in the 1970s, which were breathtaking, just breathtaking.

PHILIPS: Let me comment on that. In the paper of mine that Rob Arnott is referring to, I took the actual live track records of every domestic TAA manager (about a dozen of them, and they had 95 percent of the assets under management in TAA at the time) and performed Henriksson–Merton and Cumby–Modest tests for timing skills. I found that in the 1970s, TAA was very successful. Then, in the 1980s, the results become a little mixed. If you include the period up to and including the crash of 1987, *all* the TAA managers added value; after the crash, no one added value. But here's an interesting twist to the story: Let's say a genie came to you once a quarter or once a month, take your choice, from 1980 onwards, and whispered "buy stocks" or "buy bonds" in your ear—and the genie was never wrong. And let's say you can make the appropriate portfolio changes without transaction costs. By how much did the genie outperform a simple 60/40 mixture of stocks and bonds? It turns out that the genie's outperformance went down enormously from the precrash to the postcrash period. It dropped from about 24 percent a year to about 15 percent a year. In effect, the genie got a lot less prosperous after 1987, so it's not surprising that TAA managers found themselves in trouble.

Summary Comments

Robert Arnott
John Campbell
Peng Chen, CFA
Bradford Cornell
William Goetzmann
Brett Hammond
Campbell Harvey
Roger Ibbotson
Martin Leibowitz
Rajnish Mehra
Thomas Philips
William Reichenstein, CFA
Robert Shiller
Kevin Terhaar, CFA
Peter Williamson

MARTIN LEIBOWITZ: I think it might be interesting to just go around the table for any last comments on our topic, the equity risk premium, or for any comments on any of the papers presented today.[1]

BRETT HAMMOND: I would like to hear more discussion from Roger Ibbotson and Rob Arnott. As I have listened to the presentations today, I have been trying to decide what we could say if we were charged as a group with coming to some consensus. I'm going to assume the role of the naive observer, and in that role, I can say I have learned that in some areas, we are talking past each other and in other areas, once we clarify the definitions (or what is being measured and how), we are closer together. That understanding is useful, but what is the next step in educating our colleagues and practitioners? What would we want to tell them about their problem, which is, of course, estimating the equity risk premium looking forward? I have been wanting to ask this question all day, so now I will: What would you tell them about the equity risk premium?

[1]For Martin Leibowitz's summary of academic and practitioner research on the equity risk premium, see the Webcast of his presentation to "Research for the Practitioner: The Research Foundation Pre-Conference Workshop" held in conjunction with the AIMR 2002 Annual Conference. The Webcast is available in summer 2002 at aimr.direct.org.

ROGER IBBOTSON: What you say is to the point. First, we see a need for clarification of what we mean by the equity risk premium: I think all of us in this room see it as an expectation, not a realization; if we look at realizations, it's to help us understand expectations. But not everybody outside the room understands this distinction.

The second issue is the use of "arithmetic" versus "geometric." Every time we make a forecast, we should say whether the forecast is arithmetic or geometric and which risk-free rate we are using—U.S. T-bills, the long bond, or TIPS.

Third, we need to distinguish between yields and returns. Jeremy Siegel, for example, used realized returns, whereas others today used realized yields.

Fourth, we should always specify the forecast horizon—whether we are talking about a short or a long horizon. The risk premium for a short horizon is basically about timing, an attempt to judge whether the market is currently over- or undervalued; the risk premium for the very long horizon provides a more stable concept of what the risk premium is—namely, the long-term extra return that an investor is expected to get for taking risks, assuming the market is fairly valued.

If we could at least get these definitions delineated and clarified and let everybody know what the definitions are, it would help identify the differences among us. We are actually much more of one mind than some might think. And the theoretical analyses actually come closer to the empirical results I might have imagined before this conference.

The 4 percent (400 bps) equity risk premium forecast that I have presented here today is a geometric return in excess of the long-term government bond yield. It is a long-term forecast, under the assumption that today's market is fairly valued.

WILLIAM REICHENSTEIN: I want to make a comment in terms of asset allocation based on the geometric difference between future stock and future bond returns. Let's say that the real return on stocks is expected to be 4 percent. Of course, the numbers would depend on the assumptions used; if you use the dividend model, the real return might be 2.5 percent, and with the earnings model, it might increase to 4 percent, but in either case, we are talking about a number well below the historical 7 percent real return on stocks. If we are looking at a real return on stocks of 4 percent and a real return on bonds of 3 percent, the equity risk premium is about 1 percent, which is much lower than in the past. So, the expectation for future equity real returns is down. But for a 50/50 stock/bond portfolio, if you use the historical Ibbotson numbers of 7 percent for stocks and 2 percent for bonds, then your historical real return on a 50/50 portfolio is 4.5 percent. How much worse off are you today at an estimate of 4 percent real return on stocks and 3 percent real return on bonds? That 50/50 portfolio has 3.5 percent real return instead of 4.5 percent, and that is only a 1 percentage point difference. Part of the reason the equity risk premium is lower, it seems to me, is because the real returns on bonds are up.

ROBERT ARNOTT: That's a very good point. The 4.5 percent versus the 3.5 percent expected *portfolio* return invites the question: Why is the actuarial community allowing sponsors to use 6.5 percent as an actuarial real return assumption for their aggregate balanced pension funds? The average nominal return is 9.3 percent, and the average inflation assumption is 2.8 percent. I would say that assuming a 6.5 percent real return is irresponsible and dangerous regardless of whether the reasonable expectation for real return going forward is 4.5 percent or 3.5 percent.

KEVIN TERHAAR: I think of the risk premium as most appropriately viewed as a discount rate element corresponding to a long horizon and relative to a risk-free rate, *commensurate with the asset's risk*. The risk premium issues that we have been discussing today are not unique to the U.S. equity market. Equities or bonds, or any other asset class for that matter, should be discounted in light of the risks that the asset entails. Although there seems to be some agreement on definition and, to a lesser extent, expectations, we are still left with a question that is one step removed from the equity risk premium: What is the appropriate price of risk as we look to the future? Even if we can agree that risk is more stable and thus more easily forecastable than return, and we are able to develop agreed-upon and reasonable forward-looking risk estimates, the issue of the appropriate *price of risk* still exists. Ultimately, it is this price of risk that determines the risk premium, not only of U.S. equities, but also of any other asset class. The risk premium on the domestic equity market should not and cannot be viewed in isolation.

LEIBOWITZ: In response to Brett Hammond, I'm very impressed by the level of consensus on the view that earnings can grow only at a somewhat slower rate than GDP per capita and that no one seems to feel it can grow much more—except Roger Ibbotson, who thought EPS could grow faster than GDP because of extra earnings retention and the implicit growth estimate inherent in the high recent price-to-earnings ratio. The fact that we're basically in agreement that earnings are tightly bound to the growth in the economy has, I think, a lot of implications. Also, I think we can agree that the distinction between arithmetic and geometric is important in terms of the way these concepts are discussed and analyzed. Another important point is that the term structure that is being used to analyze the risk premium must be defined. We also need to keep in mind that the estimation error over the short term is very, very high. So, our views, at least our expectations, may be more convergent over time, but the differences still remain.

Another thing that is surprising is the disconnect between the low growth assumption and the risk premium we tend to believe in, or at least corporate executives tend to believe in. Historically, the risk premium has been more than 5 percent, which may be tough to get in the future with the earnings growth numbers that have been cited today. I think we've come to some important agreements here.

I am troubled, however, by one aspect we haven't explored: Given the growth rate of GDP (the rate of all the corporate profits—including all the entrepreneurial profits that are not captured in the public market, all the free enterprise profits in the economy), how much of the earnings has to be reinvested to sustain that growth? That's a critical equilibrium question. Roger is the only person who addressed it, which he did in terms of his historical study. I think this point is worthy of a lot more thought.

ARNOTT: In terms of the lessons learned today, a tidy way to look at the whole returns picture is to hearken back to the basic notion that the real return on stocks has just three constituent parts—changes in valuation levels, growth, and income (whether income is dividends or dividends plus buybacks). We typically know the yield, so much of the discussion gets simplified to a reexamination of two key issues: (1) Is current pricing wrong? Should valuation levels change? (2) What growth rate is reasonable to expect? As you saw in the rather sharp dichotomy between my formulation for growth and Roger Ibbotson's formulation for growth, there's plenty of room for dialogue—in fact, immense room for dialogue.

A related aspect I think is interesting to observe is that, although there are a whole host of theories relating to finance, some of them elegant, brilliantly crafted, and sensible formulations of the way the world *ought* to work—the capital asset pricing model and Modigliani and Miller being two vivid examples—comparatively few people believe that the world actually works in exact accord with any such theories. We've seen tangible evidence that M&M, while a fine theory, doesn't necessarily work intertemporally. And we know that the CAPM in its raw form doesn't fit the data very well. This doesn't make it a bad theory; it's a wonderful theory and a wonderful formulation of the way the world ought to work. Similarly, the notion that higher P/Es should, in an efficient market, imply faster future permanent growth makes sense. It's an intuitive theory. Does it stand up to historical testing? No.

A similar lesson I think we can take away from today is that the theory and the reality of the risk premium puzzle differ. There are a host of theories that relate to the risk premium puzzle and, from our views on the risk premium, relate to the asset allocation decision, but the theories don't stand up to empirical tests. A very interesting area of exploration for the years ahead will be to try to find a theoretically robust construct that fits the real world.

CAMPBELL HARVEY: I was struggling through the morning just with the vocabulary related to the risk premium: It depends on the horizon; it depends on the risk-free rate; it's a moving target through time; it's conditional; it's unconditional. I now have a better understanding of these concepts and the difficulties in defining them. It is extraordinary that, given the importance of the definitions of these variables, there is so much disagreement in terms of approach. Indeed, I have to teach this material, and it is a difficult topic for the students. We talk in class about the risk premium, but we also have to take a step back and define risk, which is extraordinarily difficult to do.

We have talked today about the current state-of-the-art models. There is a burgeoning literature on different measures of risk, and we are learning a lot from the new behavioral theories. So, we are moving forward in our understanding of the risk premium. Indeed, some of the foremost contributors to this effort are in this room. And I think more progress will be made in the future. It is somewhat frustrating that we are not there yet. I cannot go into the classroom or into the corporate world and say with some confidence, "*This* is the risk premium."

ROBERT SHILLER: I was thinking about the ambiguity of our definitions of the equity risk premium and about what we mean by expectations. We tend to blur the concepts of our own expectations with the public's expectations and with rational expectations. And the interpretations we give to the concept of expectations have changed through time. The history of thought about expectations is interesting. I remember a 1969 article by Conard and Frankena about the term structure—before the rational expectations revolution—that asserted that there is no objective way to specify expectations in a testable model but by assuming perfect foresight. They wrote this after Muth (1961) wrote the first treatise on rational expectations but before it had any impact on the profession. Without access to the theoretical framework proposed by Muth, there was no concept at all of rational expectations. That was then, and now, today, 30 years later, we economists often seem to think that the word "expectations" has no other meaning than "rational expectations."

Economists today think expectation is the summation of P_iX_i, where P is the probability, but that is a very abstract concept that we've been taught. We can trace the word "probability" very far back in time, but it didn't always have all the associations that it has today. The word "probability" didn't even have the meaning that we attach to it now until the mid-1600s, when it seemed to suddenly explode on the intellectual scene. Before then, the word "probability" existed, but it meant "trustworthiness" and had no connection at all to our modern concept of probability. Suddenly, Blaise Pascal and others got people talking about probability, which led naturally to the concept of mathematical expectation.

Just as "probability" is not a natural concept, I think "expectations" is not a natural concept. When you do surveys and you ask people for their expectations, should we expect them to give us some calculation of mathematical expectations? In fact, their reaction to questions about their expectations often seems a sort of a panic: What are these people asking for? What kind of number do they want? I have to come up with a number fast! (Incidentally, a lot of people don't remember that John Maynard Keynes' first claim to fame was a 1921 book about probability in which he argued that people really don't have probabilities as we think of them today.[2])

With all of these ambiguities, one starts to wonder what the equity risk premium is measuring. When I was surveying individual and institutional investors about their outlook for the market, I found that if I asked investors what they thought the DJIA would do in the next year, the average answer was + 5 percent. But the PaineWebber/Gallup survey taken at the same time found that investors thought the DJIA would rise by 15 percent. That's quite a big discrepancy. So, I called Gallup and asked them if we could figure out the reason for such different results. As it turned out, the different survey responses were a function of the wording of the questions. The Gallup

[2] This work can be found in Keynes (1973).

poll was conducted by randomly telephoning people at the dinner hour. Their question was (more or less): What return do you expect on the stock market in the next year in percentage terms? My survey was conducted through a written questionnaire, and the specific question about the market was (more or less): "What do you think the DJIA is going to do in the next 12 months? Put a plus mark if you think it's going to go up and put a minus mark if you think it's going to go down."

The critical difference is that I mentioned the possibility that the market might go down, so about one-third of my answers were negative. I called Gallup and asked them what fraction of their respondents said "Down." And they said that there were so few down responses that they rounded them to zero. So, I was trying to figure out why they got so few negative responses. Well, the Gallup respondents were called at dinnertime, and maybe the person who called was somewhat intimidating, so respondents had to have some courage to say they thought the market return was going to be negative. In my survey, however, I brought up in writing a possible negative choice, and I got a lot of negative responses. So, I think reported expectations are very fragile.

In the investment profession, we've learned to have respect for psychologists and the concepts they use because they've learned a lot by studying how people frame their thinking and decision making. The concepts arising from this knowledge can be very helpful to us in our work. And psychologists deal with other attitudes related to expectations—aspiration, hope, regret, fear, and the salience of stories. All of these parameters are constantly changing through time. So, when you ask someone about their expectations, the answer they give will be very context sensitive.

With surveys, we've learned you need to ask exactly the same questions in exactly the same order on each questionnaire. Even so, you don't know quite what you're really getting because expectations have so many different definitions.

RAJNISH MEHRA: I want to make two quick comments. My first point is that valuation models help us structure the problem, but what breathes life into a valuation model are the forecasts, and these forecasts have huge conditional errors. Not many of the estimates for the equity premium that were given today were accompanied by the standard deviation of that estimate. That standard deviation is too important to be missing. For example, in my data relating the expected mean equity risk premium to national income, the standard deviation around that mean is huge. Just giving a point estimate is not enough. The omission of the conditional error worries me.

My second point is that profound demographic shifts are going to be occurring in the United States, in terms of the Baby Boomers retiring, about which Ed Prescott and I wrote (1985). That phenomenon is going to lead to asset deflation, which has profound implications for the *ex ante* equity premium.

THOMAS PHILIPS: I have been very interested to see two broad strands of thought discussed today. One of these strands, exemplified by Rajnish Mehra, is the line of thinking in which the basic model involves human economic behavior, whether that behavior is utility maximizing or motivated by something else, and the effects of that behavior in the capital markets. The second strand is more empirical—constructing a point estimate for the equity risk premium—and it is exemplified by Rob Arnott's and Roger Ibbotson's work. I see two somewhat different challenges for these two strands, and ultimately, they have to meet in the middle so that we can build a unified theory.

For the economist, the challenge I see is related to Richard Feynman's argument about why scientific imagination is so beautiful: It must be consistent. You cannot imagine just anything; it has to be consistent with classical mechanics, with quantum mechanics, with general relativity, and so on and so forth. Within this set of constraints, beautiful ideas are born that tie neatly into a powerful edifice. I see the challenge for financial economists as not simply explaining the equity risk premium but explaining a fairly wide range of economic phenomena within a unified framework. Instead of a patchwork of models, financial economics needs to look more like physics.

The challenge for the second group of people, those who provide the point estimates, is (as Rajnish Mehra correctly points out) to estimate some of the errors in our estimates and to be able to communicate all this information in a language that is accessible to the person on the street. In particular, we need to dissuade investors from using the sample mean as the best estimator of the true mean.

So, the two challenges are different, but the overarching challenge is to somehow unify the two approaches in a clean way that answers the question of what the equity risk premium is and makes tactical predictions.

BRADFORD CORNELL: I like to think more in terms of valuation and expected returns than in terms of the equity risk premium. The salient feature to me in that regard is that corporate profits after tax seem to be closely tied to GNP, particularly if the market is measured properly, in the aggregate and not limited to the S&P 500 Index, so that what we have to value is not all that uncertain. However, the way we value earnings, as Rajnish Mehra pointed out, has changed quite a bit. Stock market value in the United States has varied over time from half of GNP to twice GNP, which is about where it is now. To say that earnings are twice GNP, we either have to say that the expected returns are low and are expected to remain low for the long term or that the market has simply made a mistake. The one point that I would make to practitioners, fund managers, and so forth, is that they cannot maintain a 6.5 percent actuarial assumption in light of these data.

PENG CHEN: I think there are probably two types of data: One type is what the companies and the economy reveal—the analysis that Roger Ibbotson and I are working on—and the other type is drawn from the investor's point of view—how much the investor expects from a project or a security. What I think is really interesting is that the answers are going to lie between these two dynamics. How people adjust to the dynamics, how the dynamics change people's behavior, and how that behavior affects the market are very important to observe. I think the reason we see the valuation of the market rise and fall is not necessarily because the entire investment community believes the actual risk premium has fallen or gone up or that risk rose or fell but because of this dynamic. Not all investors have to change their minds to affect market value. Maybe the dynamic affected only a small number or a certain group of investors; only a marginal number of investors have to change their minds. So, it would be interesting to see how the two sides work together dynamically.

PETER WILLIAMSON: One of the most interesting aspects of our discussion today is the areas of agreement and of disagreement. The benefit of identifying areas of disagreement is that it can lead to the search for the reason for the disagreement. It is fascinating to me how all of the findings or theory might be implemented. Can you imagine an active manager turning to his clients and saying, "You must understand that the growth in earnings of your portfolio can't exceed GDP growth"? The client wouldn't believe it, and the manager wouldn't believe it. An active manager can't afford to believe it. Or can you imagine a firm that sells S&P 500 indexed funds sending a letter to all of the shareholders saying that they must realize earnings cannot grow faster than GDP? I can't imagine that message going out. So, what impact does all of the discussion we have had today make on the actual allocation of assets, the actual management of money? I don't know. I don't know whether investors ever have to *really* understand the equity risk premium, whether it's even in their best interest to understand it.

As for allocation, my sense is that different sectors of the investment community will do very different things in terms of asset allocation on the strength of the same expected risk premium. I think that the CREF participant who's 25 years old—looking ahead 40 years to retirement, saving money—versus the investor who is 66 years old—in the process of "dissaving," consuming now—given the same expected rate of return on equity, might do very different things with their money.

Richard Thaler and I deal with the problem of college and university endowment funds. One would think that endowment funds should all be thinking very long term, but the decisions are made by people—who don't live centuries and who, in fact, can be very embarrassed if the endowment has even one very poor quarter. For example, I am on the investment committee of a prep school, and years ago, the trustees agreed that the school should be much more heavily invested in equities, that the school should be thinking long term—but not yet. And each year, the suggestion is repeated, but the decision is: not yet.

It's very, very difficult for people to think long term. Yet, to a large extent, what we've been talking about today is what's sensible for the long term. Well, if people simply cannot think long term, then we are reduced to decisions for the short term. And the asset allocation implications may be very different for investors who cannot think much beyond the next quarter from the implications for those who, in theory at least, ought to be thinking about the next 50 years.

In short, I'm really puzzled about where all that we have discussed goes in terms of making any impact on investment behavior and on asset allocation.

JOHN CAMPBELL: My starting point is that we live in a world in which the forward-looking, *ex ante* equity premium that you might expect if you're a thoughtful investor trying to be rational changes over time, and those changes have implications for the methods used to estimate the premium. We've discussed these estimation methods today, and I think we have quite a consensus that past returns can be very misleading so it is probably better to start with valuation ratios and adjust them for growth expectations.

If we live in a world in which these numbers—the real interest rate, the equity premium, and so forth—change over time, that has a big impact on asset allocation. So, I can't resist plugging my forthcoming book with Luis Viceira (2002), *Strategic Asset*

Allocation: Portfolio Choice for Long-Term Investors. Brad Cornell's colleagues at UCLA coined the term "strategic asset allocation" to contrast with tactical asset allocation (Brennan, Schwartz, and Lagnado 1997). TAA is myopic; it looks at the next period, at the risk–return in one period. The idea behind strategic asset allocation is that if risk premiums are changing over time, the risks of different asset classes may look different for different horizons. You wouldn't get such an effect if returns were identically and independently distributed, but it can become quite important if the stock market is mean reverting or if real interest rates change over time.

I'm a little more optimistic than Peter Williamson is. I think there is some hope of influencing the practical world to think about these issues, because many of the rules of thumb that financial planners have used for years have this flavor. That is, the rules make more sense in a dynamically changing world than they would in an i.i.d. world. So, there's been a mismatch between academic research and practitioners' rules of thumb. We can close that gap if we accept in our models of asset allocation that investment opportunities change over time. So, we might, with some additional work, be able to narrow the gap between how practitioners think and how academics think.

WILLIAM GOETZMANN: The thing that struck me about our discussion today is that, with the exception of Campbell Harvey's paper, almost everything we're doing is an interpretation of history—whether it's historical valuation ratios, arithmetic means, or what have you. That basis for argument is exciting but has its limitations. History, after all, is a series of accidents; the existence of the time series since 1926 might itself be an accident. So, I'm more convinced than ever that we've got to find a way out of the focus on U.S. historical data if we want to solve some of these questions and to reassure ourselves, if indeed we can, that the equity premium is of a certain magnitude.

LEIBOWITZ: Thank you all.

References and Recommended Readings

Abel, Andrew B. 1990. "Asset Pricing under Habit Formation and Catching Up with the Joneses." *American Economic Review*, vol. 80, no. 2:38–42.

Ang, Andrew, and Geert Bekaert. 2001. "Stock Return Predictability: Is it there?" National Bureau of Economic Research (NBER) Working Paper 8207 (April).

Arnott, Robert D., and Clifford S. Asness. 2002. "Does Dividend Policy Foretell Earnings Growth?" Unpublished working paper (January).

Arnott, Robert D., and Peter L. Bernstein. 2002. "What Risk Premium is 'Normal'?" *Financial Analysts Journal*, vol. 58, no. 2 (March/April 2002):64–85.

Arnott, Robert D., and James N. von Germeten. 1983. "Systematic Asset Allocation." *Financial Analysts Journal,* vol. 39, no. 6 (November/December):31–38.

Arnott, Robert D., and Ronald Ryan. 2001. "The Death of the Risk Premium: Consequences of the 1990s." *Journal of Portfolio Management*, vol. 27, no. 3 (Spring):61–74.

Asness, Clifford S. 2000a. "Bubble Logic or How to Learn to Stop Worrying and Love the Bull." AQR working paper (June).

———. 2000b. "Stocks versus Bonds: Explaining the Equity Risk Premium." *Financial Analysts Journal*, vol. 56, no. 2 (March/April):96–113.

Bagwell, Laurie Simon, and John B. Shoven. 1989. "Cash Distributions to Shareholders." *Journal of Economic Perspectives*, vol. 3, no. 3 (Summer):129–140.

Bansal, Ravi, and Amir Yaron. 2000. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." NBER Working Paper 8059 (December).

Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad. 2001. "Consumption, Dividends, and the Cross-Section of Equity Returns." Working paper, Duke University (October).

Barberis, N., M. Huang, and T. Santos. 1996. "Prospect Theory and Asset Prices." Working paper, University of Chicago.

Benartzi, S., and R. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics*, vol. 110, no. 1 (February):73–92.

Bernstein, Peter L. 1993. *Volatility: Two Views, Two Signals*. Newsletter. (November 15).

———. 1996. "Are Stocks the Best Place to Be in the Long Run? A Contrary Opinion." *Journal of Investing*, vol. 5, no. 2 (Summer):6–9.

———. 1997. *Stock/Bond Risk Perceptions and Expected Returns*. Newsletter. (February 1).

———. 2001a. "Ending Endism." *Economics and Portfolio Strategy* (April 30).

———. 2001b. "The Real Economy and the Unreal Economy." *Economics and Portfolio Strategy* (May 15).

Black, Fisher. 1976. "Studies of Stock Price Volatility Changes." *Proceedings of the 1976 Meetings of the American Statistical Society, Business and Economics Statistics Section*:177–181.

Bogle, John C. 1991. "Investing in the 1990s: Occam's Razor Revisited." *Journal of Portfolio Management*, vol. 18, no. 1 (Fall):88–91.

———. 1995. "The 1990s at the Halfway Mark." *Journal of Portfolio Management*, vol. 21, no. 4 (Summer):21–31.

Boudoukh, Jacob, and Matthew Richardson. 1994. "The Statistics of Long-Horizon Regressions Revisited." *Mathematical Finance*, vol. 4, no. 2 (April): 103–119.

Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado. 1997. "Strategic Asset Allocation." *Journal of Economic Dynamics and Control,* vol. 21, no. 8–9 (June):1377–1403.

Campbell, J.Y. Forthcoming 2002. "Consumption Based Asset Pricing." In *Handbook of the Economics of Finance*. Edited by G.M. Constantinides, M. Harris, and R. Stulz. Amsterdam, Netherlands: North-Holland.

Campbell, John Y., and J.H. Cochrane. 1995. "By Force of Habit: A Consumption-Based Explanation of

Aggregate Stock Market Behavior." NBER Working Paper 4995.

———. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy*, vol. 107 (April):205–251.

Campbell, John Y., and Robert J. Shiller. 1988. "Stock Prices, Earnings, and Expected Dividends." *Journal of Finance*, vol. 43, no. 3 (July):661–676.

———. 1998. "Valuation Ratios and the Long-Run Stock Market Outlook." *Journal of Portfolio Management*, vol. 24, no. 2 (Winter):11–26.

———. 2001. "Valuation Ratios and the Long-Run Stock Market Outlook: An Update." NBER Working Paper 8221 (April).

Campbell, John Y., and Luis M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford, U.K.: Oxford University Press.

Chancellor, Edward. 1999. *Devil Take the Hindmost*. New York: Farrar, Straus, and Giroux.

Claus, James, and Jacob Thomas. 2001. "Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stocks." *Journal of Finance*, vol. 56, no. 5 (October):1629–66.

Cole, Kevin, Jean Helwege, and David Laster. 1996. "Stock Market Valuation Indicators: Is This Time Different?" *Financial Analysts Journal*, vol. 52, no. 3 (May/June):56–64.

Conard, Joseph W., and Mark W. Frankena. 1969. In *Essays on Interest Rates*. Edited by Jack M. Guttentag and Philip Cagan. New York: NBER and Columbia University Press.

Constantinides, George C. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy*, vol. 98 (June):519–543.

Constantinides, George M., John B. Donaldson, and Rajnish Mehra. 1998. "Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle." NBER Working Paper 6617 (June).

Cornell, Bradford. 1999. *The Equity Risk Premium*. New York: John Wiley & Sons.

Cowles, Alfred. 1938 (1939, 2nd edition). *Common Stock Indexes: 1871–1937*. Bloomington, IN: Principia.

Cragg, John G., and Burton G. Malkiel. 1982. *Expectations and the Structure of Share Prices*. Chicago, IL: University of Chicago Press.

Damodaran, Aswath. Forthcoming 2002. "Choosing the Right Valuation Model." In *Equity Valuation Techniques*. Charlottesville, VA: AIMR.

Davidson, L.B., and D.O. Cooper. 1976. "A Simple Way of Developing a Probability Distribution of Present Value." *Journal of Petroleum Technology* (September):1069–78.

Davis, S.J., and Rajnish Mehra. Forthcoming 2002. *Stocks and Bonds in the Portfolio Life-Cycle*. Saddle River, NJ: Prentice Hall.

DeLong, B., A. Shleifer, L. Summers, and R. Waldmann. 1990. "Noise Trader Risk in Financial Markets." *Journal of Political Economy*, vol. 98, no. 4 (August):703–738.

Diermeier, Jeffrey, Roger Ibbotson, and Laurence B. Siegel. 1984. "The Supply of Capital Market Returns." *Financial Analysts Journal*, vol. 40, no. 2 (March/April):74–80.

Dimson, Elroy, Paul Marsh, and Mike Staunton. 2001. *Triumph of the Optimists: 101 Years of Global Investment Returns*. Princeton, NJ: Princeton University Press.

Epstein, L., and S. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, vol. 57, no. 4 (July):937–969.

Fama, Eugene F., and Kenneth R. French. 1988. "Dividend Yields and Expected Stock Returns." *Journal of Financial Economics*, vol. 22, no. 1 (October):3–26.

———. 1989. "Business Conditions and Expected Returns on Stocks and Bonds." *Journal of Financial Economics*, vol. 25, no. 1 (July):23–50.

———. 2000. *The Equity Risk Premium*. Center for Research in Security Prices Working Paper 522 (April).

Fama, Eugene F., and Kenneth R. French. 2001. "The Equity Premium." Center for Research in Security Prices Working Paper No. 522 (April).

Fisher, Irving. 1930. *The Theory of Interest*. New York: MacMillan.

Frankel, Jeffery, and Kenneth Froot. 1990. "Chartists, Fundamentalists, and Trading in the Foreign

Exchange Market." *American Economic Review*, vol. 80, no. 2 (May):181–185.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh. 1987. "Expected Stock Returns and Volatility." *Journal of Financial Economics*, vol. 19, no. 1 (January):3–29.

Garber, Peter M. 2000. *Famous First Bubbles: The Fundamentals of Early Manias*. Cambridge, MA: MIT Press.

Glassman, James K., and Kevin A. Hassett. 1999. *Dow 36,000*. New York: Random House.

Goetzmann, William N., and Philippe Jorion. 1993. "Testing the Predictive Power of Dividend Yields." *Journal of Finance*, vol. 48, no. 2 (June):663–679.

Goetzmann, William N., and Massimo Massa. 2001. "Dispersion of Opinion and Stock Returns: Evidence from Index Fund Investors." Yale ICF Working Paper No. 00-55 (August).

Goyal, Amit, and Ivo Welch. 1999. "Predicting the Equity Premium." Working paper, University of California at Los Angeles.

———. 2001. "Predicting the Equity Premium with Dividend Ratios." Yale ICF working paper, University of California at Los Angeles (December).

Graham, B., and D.L. Dodd. 1934. *Security Analysis*. New York: McGraw-Hill.

Graham, John, and Campbell R. Harvey. 1996. "Market Timing Ability and Volatility Implied in Investment Newsletters' Asset Allocation Recommendations." *Journal of Financial Economics*, vol. 42 (January):397–422.

———. 2001a. "Expectations of Equity Risk Premia, Volatility and Asymmetry from a Corporate Finance Perspective." Working paper, Duke University.

———. 2001b. "The Theory and Practice of Corporate Finance: Evidence from the Field." *Journal of Financial Economics*, vol. 60, no. 2–3 (May):187–243.

———. 2002. "How Do CFOs Make Capital Budgeting and Capital Structure Decisions?" Working paper, Duke University (March).

Granger, C.J.W., and P. Newbold. 1974. "Spurious Regressions in Econometrics." *Journal of Econometrics*, vol. 2:111–120.

Hicks, Sir John Richard. 1946. *Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory*. Oxford, UK: Oxford University Press.

Ibbotson, Roger, and Peng Chen. 2002. "Stock Market Returns in the Long Run: Participating in the Real Economy." Yale ICF Working Paper (March).

Ibbotson, Roger, and Rex A. Sinquefield. 1976a. "Stocks, Bonds, Bills and Inflation: Year-by-Year Historical Returns (1926–74)." *Journal of Business*, vol. 49, no. 1 (January):11–47.

———. 1976b. "Stocks, Bonds, Bills and Inflation: Simulations of the Future (1976–2000)." *Journal of Business*, vol. 49, no. 3 (July):313–338.

Ibbotson Associates. 2001. *Stocks, Bonds, Bills, and Inflation: 2000 Yearbook*. Chicago, IL: Ibbotson Associates.

Ilmanen, Antti. 1995. "Time-Varying Expected Returns in International Bond Markets." *Journal of Finance*, vol. 50, no. 2 (June):481–506.

Jagannathan, Murali, Clifford P. Stephens, and Michael S. Weisbach. 2000. "Financial Flexibility and the Choice between Dividends and Stock Repurchases." *Journal of Financial Economics*, vol. 57, no. 3 (September):355–384.

Jagannathan, Ravi, and Iwan Meier. 2001. "Do We Need the CAPM for Capital Budgeting?" Working paper, Northwestern University.

Jagannathan, Ravi, Ellen R. McGrattan, and Anna Scherbina. 2000. "The Declining U.S. Equity Premium." *Quarterly Review*, Federal Reserve Bank of Minneapolis, vol. 24, no. 4 (Fall):3–19.

Janis, Irving. 1982. *Groupthink*. Boston, MA: Houghton Mifflin.

Jones, Charles. 2001. "A Century of Stock Market Liquidity and Trading Costs." Paper presented to the NBER Summer Institute, Asset Pricing Meetings.

Kahneman, Daniel, and Amos Tversky. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science*, vol. 185, no. 4157 (September):1124–31.

———. 1979 (and 1974). "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, vol. 46, no. 2 (March):171–185.

Kane, Alex, Alan J. Marcus, and Jaesun Noh. 1996. "The P/E Multiple and Market Volatility." *Financial Analysts Journal*, vol. 52, no. 4 (July/August):16–24.

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest, and Money*. New York: Harcourt Brace and World.

———. 1973. *Treatise on Probability* [1921]. In *The Collected Writings of John Maynard Keynes*. London, U.K.: MacMillan, for the Royal Economic Society.

Kirby, Chris. 1997. "Testing the Predictable Variation in Stock and Bond Returns." *Review of Financial Studies*, vol. 10, no. 3 (Fall):579–630.

Knight, Frank. 1964. *Risk, Uncertainty and Profit*. New York: Augustus M. Kelley.

Lander, Joel, Athanasios Orphanides, and Martha Douvogiannis. 1997. "Earnings Forecasts and the Predictability of Stock Returns: Evidence from Trading the S&P." *Journal of Portfolio Management*, vol. 23, no. 4 (Summer):24–35.

Lintner, John. 1956. "Distribution of Incomes of Corporations among Dividends, Retained Earnings, and Taxes." *American Economic Review*, vol. 46:97–113.

Macaulay, Frederick R. 1938. *Some Theoretical Problems Suggested by the Movements of Interest Rates, Bond Yields and Stock Prices in the United States since 1856*. Washington, DC: National Bureau of Economic Research.

McDonald, R., and Dan Siegal. 1986. "The Value of Waiting to Invest." *Quarterly Journal of Economics*, vol. 101, no. 4:707–728.

McGrattan, Ellen R., and Edward C. Prescott. 2000. "Is the Stock Market Overvalued?" *Federal Reserve Bank of Minneapolis Quarterly Review*, vol. 24, no. 4 (Fall):20–40.

McGrattan, Ellen R., and Edward C. Prescott. 2001. "Taxes, Regulations, and Asset Prices." Unpublished paper, Federal Reserve Bank of Minneapolis (March).

Mehra, Rajnish. 1998. "On the Volatility of Stock Prices: An Exercise in Quantitative Theory." *International Journal of Systems Science*, vol. 29, no. 11:1203–11.

———. Forthcoming 2002a "Equity Premium Puzzle." In *Mastering Investments*. Edited by James Pickford. Saddle River, NJ: Prentice Hall.

———. Forthcoming 2002b. "The Equity Premium: Why Is It a Puzzle?" *Financial Analysts Journal*.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economies*, vol. 15, no. 2 (March):145–161.

———. 1988. "The Equity Risk Premium: A Solution?" *Journal of Monetary Economics*, vol. 22 (July):133–136.

———. Forthcoming 2002. "The Equity Premium Puzzle in Retrospect." In *Handbook of the Economics of Finance*. Edited by G.M. Constantinides, M. Harris, and R. Stulz. Amsterdam, Netherlands: North Holland.

Merton, R.C. 1980. "On Estimating the Expected Return on the Market." *Journal of Financial Economics*, vol. 8, no. 4 (June):323–361.

Miller, Merton, and Franco Modigliani. 1961. "Dividend Policy, Growth, and the Valuation of Shares." *Journal of Business*, vol. 34, no. 4 (October):411–433.

Modigliani, F., and R. Cohn. 1979. "Inflation, Rational Valuation, and the Market." *Financial Analysts Journal*, vol. 35, no. 2 (March/April):24–44.

Modigliani, Franco, and Merton H. Miller. 1958. "The Cost of Capital, Corporation Finance, and the Theory of Investment." *American Economic Review*, vol. 48, no. 3 (June):261–297.

Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica*, vol. 29:315–335.

Philips, Thomas K. 1999. "Why Do Valuation Ratios Forecast Long-Run Equity Returns?" *Journal of Portfolio Management*, vol. 25, no. 3 (Spring):39–44.

Philips, T.K., G.T. Rogers, and R.E. Capaldi. 1996. "Tactical Asset Allocation: 1977–1994." *Journal of Portfolio Management*, vol. 23, no. 1 (Fall):57–64.

Poterba, James M., and Lawrence H. Summers. 1986. "The Persistence of Volatility and Stock Market Fluctuations." *American Economic Review*, vol. 76, no. 5 (December):1142–51.

Reichenstein, William. Forthcoming 2002. "What do Past Stock Market Returns Tell Us about the Future?" *Journal of Financial Planning*.

Samuelson, Paul A. 1994. "The Long-Term Case for Equities and How It Can Be Oversold." *Journal of Portfolio Management*, vol. 21, no. 1 (Fall):15–26.

———. 1979. "Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long." *Journal of Banking and Finance*, vol. 3, no. 4:305–308.

Schwert, G. William. 1989. "Why Does Stock Market Volatility Change over Time?" *Journal of Finance*, vol. 44, no. 5 (December):1115–54.

Schwert, G. William. 1990. "Indexes of United States Stock Prices from 1802 to 1987." *Journal of Business*, vol. 63, no. 3 (July):399–426.

Shafir, E., I. Simonson, and A. Tversky. 2000. "Reason-Based Choice." In *Choices, Values, and Frames*. Edited by D. Kahneman and A. Tversky. Cambridge, U.K.: Cambridge University Press.

Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review*, vol. 73, no. 3 (June):421–436.

———. 1989. *Market Volatility*. Boston, MA: Massachusetts Institute of Technology Press.

———. 1990. "Market Volatility and Investor Behavior." *American Economic Review*, vol. 80:58–62.

———. 2000. *Irrational Exuberance*. Princeton, NJ: Princeton University Press.

———. 2002. "Bubbles, Human Judgment, and Expert Opinion." *Financial Analysts Journal*, vol. 58, no. 3 (May/June):18–26.

Shleifer, Andrei, and Nicholas Barberis. 2000. "Style Investing." NBER Working Paper No. 8039 (December).

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

———. 1999. "The Shrinking Equity Premium." *Journal of Portfolio Management*, vol. 26, no. 1 (Fall):10–17.

———. 2002. *Stocks for the Long Run*. 3rd ed. New York: McGraw Hill.

Smith, Edgar Lawrence. 1924. *Common Stocks as Long-Term Investments*. New York: MacMillan.

Sorensen, Eric H., and Robert D. Arnott. 1988. "The Risk Premium and Stock Market Performance." *Journal of Portfolio Management*, vol. 14, no. 4 (Summer):50–55.

Stambaugh, Robert F. 1984. "Bias in Regressions with Lagged Stochastic Regressors." Manuscript.

Statman, Meir. 2002. "Lottery Players/Stock Traders." *Financial Analysts Journal*, vol. 58, no. 1 (January/February):14–21.

Thaler, R., A. Tversky, D. Kahneman, and A. Schwartz. 1997. "The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test." *Quarterly Journal of Economics*, vol. 112, no. 2 (May):647–661.

Welch, Ivo. 2000. "Views of Financial Economists on the Equity Premium and on Financial Controversies." *Journal of Business*, vol. 73, no. 4 (October):501–537.

Wilson, Jack W., and Charles P. Jones. Forthcoming 2002. "An Analysis of the S&P 500 Index and Cowles Extensions: Price Indexes and Stock Returns, 1870-1999." *Journal of Business*.

Note on Value Drivers[1]

Value-based management assumes that value creation should be a primary consideration in managerial decision making. It requires a thorough understanding of what creates value and why as well as the ability to measure value accurately. The goal of this note is to highlight the determinants of equity value and, in doing so, provide a framework for making financial, strategic, and investment decisions. In particular, the note describes three value drivers: profitability, advantage horizon, and reinvestment. Using both a theoretical model and a numerical example, it shows how each value driver affects equity value and explains why. It also presents empirical evidence to support the relation between the value drivers and value creation.

Theoretical Equity Valuation Model

Discounted cash flow (DCF) analysis translates future cash flows into current market values. For example, given a stream of equity cash flows (ECF) and a discount rate equal to the cost of equity (K_E), the market value of equity (E_{MV}) is the present value of future equity cash flows:

$$E_{MV} = ECF_1/(1+K_E) + ECF_2/(1+K_E)^2 + \ldots \qquad (1)$$

When the equity cash flows and discount rate are constant over time, this series is a stable perpetuity which can be written as:

$$E_{MV} = ECF/K_E \qquad (2)$$

Assuming that the equity cash flows are equal to the accounting return on equity (ROE) times the book value of equity (E_{BV}) at the beginning of the period, then equation 2 can be rewritten as:

$$E_{MV} = [(ROE)^*(E_{BV})]/K_E \qquad (3)$$

$$\text{where } ROE = \text{Net Income}/E_{BV}$$

While the assumption that equity cash flows are equal to accounting earnings is convenient for expositional reasons, this assumption is clearly not valid except in very special circumstances. For example, non-cash items such as depreciation or deferred taxes, and cash-items that do not flow through the income statement such as changes in working capital and fixed assets both cause cash

[1] Much of the material in this note appears in Fruhan (1979), chapter 1.

Professor Benjamin C. Esty prepared this note as the basis for class discussion.

flows to deviate from reported net income. Nevertheless, this assumption is not a bad approximation and, as will be shown in the next section, seems to generate reasonable empirical predictions.

After dividing each side of equation 3 by the book value of equity, the left side of the equality becomes the market-to-book ratio (the market value of equity divided by the book value of equity):

$$\text{Market/Book} = E_{MV} / E_{BV} = \text{ROE}/K_E \quad (4)$$

Equation 4 says that a firm's market-to-book ratio equals the ratio of its return on equity to its cost of equity. This simple valuation model, or variations of it, can be used to analyze the relation between profitability, growth, and value.

Profitability

The first value driver, profitability, is immediately clear from equation 4. For a given industry, more profitable firms—those able to generate higher returns per dollar of equity—should have higher market-to-book ratios. Conversely, firms which are unable to generate returns in excess of their cost of equity should sell for less than book value.

Profitability	Value
If ROE > K_E	then Market/Book > 1
If ROE = K_E	then Market/Book = 1
If ROE < K_E	then Market/Book < 1

One implication of this model is that firms can increase equity value by increasing their return on equity. The Du Pont formula decomposes ROE into three components and provides some guidance on how to increase it:

$$\begin{aligned} \text{ROE} &= (\text{Net Income/Equity}) \\ &= (\text{Net Income/Sales}) * (\text{Sales/Assets}) * (\text{Assets/Equity}) \\ &= (\text{Profit Margin}) * (\text{Asset Turnover}) * (\text{Financial Leverage}) \end{aligned}$$

For example, increasing the profit margin through higher prices or lower costs will increase the ROE. Similarly, increasing the asset turnover by increasing inventory turnover or reducing days receivables will increase the ROE. However, increasing financial leverage has dual, and possibly contradictory, effects. It increases not only the ROE through the Du Pont formula, but also the cost of equity.

A firm's cost of equity, or equivalently investors' expected return on equity, can be estimated using the Capital Asset Pricing Model (CAPM). According to the model, the expected return on equity is a function of a firm's equity beta (β_E) which, in turn, is a function of both leverage and asset risk (β_A):

$$K_E = R_F + \beta_E (R_M - R_F) \quad (5)$$

where:

R_M = return on the market portfolio
R_F = risk-free rate of return

$$\beta_E = [\, \beta_A - \beta_D (D/V) \,] (V/E) \quad (6)$$

because:

$$\beta_A = \beta_D (D/V) + \beta_E (E/V) \quad (7)$$

and

$$\text{Firm Value (V)} = \text{Debt Value (D)} + \text{Equity Value (E)} \quad (8)$$

Assuming riskless debt, meaning the beta of debt is zero, then equation 6 can be written as:

$$\beta_E = \beta_A \ (V/E) \quad (9)$$

As financial leverage (D/V) increases, the ratio of firm value to equity value (V/E) increases, the equity beta increases, and, according to equation 5, the expected return on equity increases. The expected return increases because equity cash flows are riskier: leverage increases debtholders fractional claim on the firm's cash flows. As a result, an increase in leverage can either increase or decrease the ratio in equation 4 depending on whether the return on equity (the numerator) or the cost of equity (the denominator) increases faster.

Advantage Horizon

Equation 4 presents a firm's market-to-book ratio as a stable perpetuity under the assumption that its profitability remains constant forever. An alternative, and more realistic assumption, is that firms generate positive abnormal returns—returns in excess of their cost of capital—for only a limited number of years. The period during which firms generate positive abnormal returns is known as the advantage horizon.

Using a variation of the simple valuation model in equation 4, Appendix 1 derives the market-to-book ratio as an annuity rather than a stable perpetuity. It assumes that a firm's equity returns can be divided into two parts: *normal* returns equal to the firm's cost of equity (K_E) and *abnormal* returns equal to the actual ROE less the cost of equity (ROE - K_E). Viewed in this fashion, one can think of abnormal returns and the advantage horizon in the same way Stewart (1991) defines economic value added (EVA) and the competitive advantage period (CAP). Equation A1.8 from the Appendix 1 is:[2]

$$\textbf{Market/Book} = 1 + (ROE-K_E) * [(1/K_E) - (1/(K_E(1+K_E)^n)] \quad (10)$$

where the advantage horizon is defined as n years. According to this formula, the greater the spread between a firm's return on equity and its cost of equity (ROE - K_E), the longer the advantage horizon (increasing n), and the sooner abnormal returns occur (positive abnormal returns in early years), the higher the market-to-book ratio. Firms that earn normal returns (K_E = ROE) in all periods should have market-to-book ratios equal to one; firms that generate negative abnormal returns during the advantage (disadvantage) period should have market-to-book ratios less than one.

Equation 10 is more realistic than equation 4 because most firms earn positive abnormal returns for only a limited number of years. The presence of positive abnormal returns encourages entry by new firms and increased competition by existing firms. Over time, competition reduces excess returns to the point where firms just earn the expected, or normal, rate of return. Although there is typically an inverse relation between the magnitude of positive abnormal profits and the length of the advantage horizon, this model implies that firms should seek to extend the advantage horizon as long as possible for a given level of profitability.

Ghemawat (1991) refers to this ability to preserve competitive advantage as sustainability and asserts it is a key determinant of value creation. Sustainability, he maintains, depends on a firm's ability to create scarcity value and for the firm's owners to capture or appropriate this value. Threats to scarcity value include imitation and substitution. A firm can defend against imitation by erecting barriers to entry or forestalling entry through aggressive positioning; a firm can defend against substitution by continually improving or augmenting its product. Threats to appropriability include

[2] This formula is a variation of the accounting-based valuation methods described in Bernard (1994); Palepu, Bernard, and Healy (1996), and Ohlson (1995).

slack and hold-up both of which result from misaligned incentives. Slack occurs when firms fail to create as much value as they are capable of creating; hold-up occurs when non-owners, instead of owners, capture value. Non-owners are often able to capture value when they provide complementary, and necessary, inputs.

Reinvestment

The third value driver, reinvestment, builds on the other two factors and incorporates the concept of growth. Firms that have attractive investment opportunities, meaning that investments are expected to generate positive abnormal earnings, can create equity value by reinvesting earnings or by investing additional equity. Appendix 2 derives a valuation model which allows for reinvestment of earnings at rate γ where γ equals the retention rate or the fraction of net income reinvested in the firm. The quantity γROE is a firm's sustainable growth rate, the rate at which it can grow its assets (or sales if they are proportional to assets) without changing its capital structure or raising external equity. With reinvestment, the valuation model becomes (equation A2.4):

$$\textbf{Market/Book} = [\textbf{ROE}(1 - \gamma)] / (\textbf{K}_E - \gamma\textbf{ROE}) \qquad (11)$$

When a firm pays out all of its earnings as dividends, then the retention rate is zero ($\gamma = 0$) and equation 11 reduces to the simple valuation model in equation 4. Assuming a firm has attractive investment opportunities in which it can generate positive abnormal returns ($ROE>K_E$), then it can increase value by retaining a larger fraction of earnings and investing them in the business. Thus reinvestment and growth creates value only when a firm can generate positive abnormal returns on future investment opportunities. Those firms with the greatest number and the most profitable investment opportunities should have the highest market-to-book ratios provided they are able to fund the projects.

In fact, it is often convenient to think of firm value as consisting of two parts: the present value of assets in place and the present value of future growth opportunities (Myers, 1977). The former require little in the way of additional investment, while the latter are investment opportunities which are expected to earn positive abnormal returns. These investment opportunities are called "real" options because they resemble financial options, particularly call options. They can be interpreted and managed using option pricing theory and valued using option pricing techniques (see Luehrman, 1995).

Numerical Example

Combining equations 10 and 11 produces a single valuation model that incorporates all three value drivers. Exhibit 1 shows this model as well as the relation between a hypothetical firm's market-to-book ratio and the value drivers. The exhibit presents three cases with differing levels of reinvestment (γ = 0%, 33%, and 66%). For each case, there is a sensitivity table showing how the market-to-book ratio depends on the advantage horizon and level of profitability (ROE).

Case #1 (no reinvestment) shows that more profitable firms have higher market-to-book ratios—the ratio increases as one reads across the rows. As stated earlier, the impact of the advantage horizon depends on whether a firm generates positive or negative abnormal earnings. The longer a firm can generate positive abnormal earnings, the greater its market-to-book ratio. However, because of discounting, abnormal earnings in later years have a smaller impact on the market-to-book ratio than abnormal earnings in early years. Alternatively, firms that generate negative abnormal earnings have market-to-book ratios less than one. Moreover, their market-to-book ratio falls as the advantage (disadvantage) horizon gets longer. Finally, the market-to-book ratio is equal to one and is independent of the advantage horizon for firms that generate normal earnings (the case where $ROE=K_E$).

Cases #2 and #3 (with reinvestment rates equal to 33% and 66%, respectively) illustrate the impact of reinvestment. Like the advantage horizon, reinvestment creates additional value only for firms that generate positive abnormal earnings. When firms are able to generate positive abnormal returns (ROE = 25%), have a long advantage horizon (30 years), and reinvest a large fraction of earnings ($\gamma = 66\%$), they create significant value. The difference between the market-to-book ratio in the high return/long horizon with no reinvestment (case #1) and with reinvestment (case #3) is large: 1.66 vs. 4.27.

Empirical Evidence

This section presents empirical evidence on the relation between the value drivers and value creation. Despite the assumptions imbedded in the simple valuation models, they do, nonetheless, yield predictions which are consistent with what we observe in practice.

Profitability

The model predicts that there is a relation between a firm's market-to-book ratio and the ratio of its return on equity to its cost of equity. Given a set of firms in a single industry, the model implies that there should be a positive relation between ROE's and market-to-book ratios for these firms assuming their costs of capital are approximately equal. To a first approximation, it is reasonable to assume that firms in the same industry will have similar capital costs because they hold similar assets and, typically, have similar capital structures.

Exhibit 2 shows the relation between market-to-book ratios and firm profitability for two quite different industries: grocery stores and oil field service companies. Whereas the grocery industry is a retail business with high inventories and low margins, the oil-field services industry is a service business with industrial customers and higher margins. Yet in both cases, there is a very clear, positive relation between equity value and ROE's: higher ROE's are associated with higher market-to-book ratios. Fruhan (1996) presents similar evidence for a much wider range of industries including newspapers, telecommunications, and specialty chemicals.

There are at least two reasons why this relation does not hold perfectly. First, not all firms in the same industry have the same leverage or same asset risk. Thus, financial and operating differences cause the cost of equity to differ across firms. Second, accounting data is subject to manipulation by managers. On the one hand, managers provide valuable information through their choice of accounting disclosures and policies. On the other hand, they are biased which may lead them to distort reported numbers. Fortunately, however, most distortions occur through accruals which eventually get reversed. Because accounting data is subject to this kind of manipulation, it is critical to understand whether the reported numbers reflect economic reality. To the extent high ROE's reflect economic reality, and not unreasonable deferral of costs or a one-time aberrations, then the relation shown in exhibit 2 will be stronger. When accounting data does not reflect economic reality, one must undo the distortions before trying to make substantive conclusions about the business or its prospects.

Advantage Horizon

Several researchers have studied the length of the advantage horizon. For example, Fruhan (1995) examined a sample of 87 "high-performing" firms defined as those firms with sales of greater than $200 million and an average ROE of greater than 25% for five consecutive years between 1976-82. He calculated the median ROE for the firms from 1976-78 and from 1989-93, and then compared these medians against the average ROE for firms on the S&P 400 (see Exhibit 3). Whereas the median ROE for the high-performing subgroup was 21% above the average ROE for the S&P 400 in 1976-82, it was only 2% above in the later period. Thus the high-performing firms' abnormal earnings had largely dissipated over the fifteen year interval.

Palepu *et al* (1996, pp. 5.4-5.7) report similar findings: abnormally high or low ROE's tend to revert to normal levels, roughly between 10-14%, often within five years and usually within ten years.[3] The reversion in ROE's is largely due to reversion in profit margins rather than reversion in asset turnover or leverage which remain relatively constant over time. The fact that advantage horizon lasts for five or ten years provides some justification for using five or ten-year projections in discounted cash flow analysis.

In another study, Ghemawat (1991) examined the returns on investment (ROI) for 692 business units from 1971-1980. After sorting the business units by their ROI in 1971, he divided the sample into two equal subgroups and calculated the average ROI for each subgroup over the next ten years. Initially, the top group had an average ROI of 39% compared to 3% for the bottom group. The 36% spread between the two groups decreased to less than 3% by the end of ten years: the average ROI for the top group had decreased to 21.5% while the average ROI for the bottom group increased to 18.0%.

While the evidence consistently shows that the advantage horizon is finite, firms like Coca-Cola, Wal-Mart, and Microsoft have been able to extend their advantage horizons for many years. These firms have been able to create tremendous value for shareholders by sustaining their ability to generate positive abnormal profits.

Reinvestment

The key insight from the model regarding investment is that reinvestment of earnings is value enhancing only when investment opportunities generate expected returns in excess of the cost of equity (ROE>K_E). Because investment opportunities vary across firms and vary over time for the same firm, it is impossible to make conclusive statements on the value of reinvestment. Nevertheless, there is some evidence that reinvestment creates value. Recent studies have shown that firms which announce major capital expenditure or research and development (R&D) programs experience positive abnormal equity returns.[4] The market interprets these announcements as good news and their stock prices usually increase. While it may be the case that firms announce only their most positive NPV investments, Fruhan (1979, Table 1-6) provides evidence from a sample of almost 1500 firms that broadly supports the relation among high profitability, high reinvestment, and high equity valuations.

Acquisitions represent another form of investment for many firms. Jensen and Ruback (1983) review the many studies on acquirer returns surrounding merger announcements. They conclude that, on average, acquirer shareholders do not lose and target shareholders gain from merger

[3] See also Freeman, Ohlson, and Penman (1982).
[4] McConnell and Muscarella (1985) analyze capital expenditure announcements while Chan, Martin, and Kensigner (1990) analyze R&D expenditure announcements.

announcements. Thus, acquisitions create net gains for both firms combined even though they do not increase acquirer shareholder value.

Jensen (1986, 1993) presents an opposing view. He argues that managers often overinvest, i.e. invest in negative net present value projects, especially when their firms generate substantial free cash flow. Their incentive to overinvest results from their compensation being tied, indirectly, to firm size which, in turn, is a function of the amount investment. They are able to over invest because internal control systems such as board oversight are weak. In the absence of effective internal control systems, external forces such as the market for corporate control discipline investment activity. Jensen cites the oil industry in general and the Gulf Oil takeover in particular as examples where takeovers eliminated wasteful capital expenditures. Just as investing in positive NPV projects creates value, so, too, does eliminating negative NPV investments.

Warren Buffet, the prominent investor and chairman of Berkshire Hathaway, acknowledged the problem of overinvestment in his company's 1984 annual report:

> Many corporations that show consistently good returns have, indeed, employed a large portion of their retained earnings on an economically unattractive, even disastrous, basis. Their marvelous core businesses camouflage repeated failures in capital allocation elsewhere (usually involving high-priced acquisitions). The managers at fault periodically report on the lessons they have learned from the latest disappointment. They then usually seek out future lessons. (Failure seems to go to their heads.). . . In such cases, shareholders would be far better off if the earnings were retained to expand only the high-return business, with the balance being paid in dividends or used to repurchase stock...

Although stated in his characteristically droll way, Buffet's point is clear: reinvestment destroys value unless it generates an appropriate risk-adjusted rate of return.

References

Bernard, Victor L, 1994, Accounting-Based valuation methods, Determinants of market-to-Book Ratios, and Implications for Financial Statement Analysis," University of Michigan working paper.

Chan, Su Han, John D. Martin, and John W. Kensigner, 1990, Corporate research and development expenditures and share value, *Journal of Financial Economics* , 26: 255-276.

Freeman, Robert, James Ohlson, and Stephen Penman, 1982, Book Rate-of-Return and Prediction of Earnings Changes: An Empirical Investigation, *Journal of Accounting Research* , Autumn: 639-653.

Fruhan, William F., 1996, How do we link together profitability, cash flows, capital cost, growth, and stock value?, Harvard Business School presentation.

Fruhan, William F., 1979, Financial Strategy: Studies in the Creation, Transfer, and Destruction of Shareholder Value, (Irwin, Homewood, IL).

Ghemawat, Pankaj, 1991, Commitment, (The Free Press, Macmillan, Inc., New York, NY).

Jensen, Michael C., 1986, Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers, *American Economic Review*, 2:323-329.

Jensen, Michael C., 1993, The Modern Industrial Revolution, Exit, and the Failure of Internal Control Systems, *Journal of Finance*, 3:831-880.

Jensen, Michael C., and Richard S. Ruback, 1983, The market for Corporate Control: The Scientific Evidence, *Journal of Financial Economics*, 11:5-50.

Luehrman, Timothy A., Capital Projects as Real Options: An introduction, Harvard Business School Note N9-295-074 (Rev. March 22, 1995)

McConnell, John J., and Chris J. Muscarella, 1985, Corporate Capital Expenditure Decisions and the Market Value of the Firm, *Journal of Financial Economics*, 14: 399-422.

Myers, Stewart C., 1977, Determinants of Corporate Borrowing, *Journal of Financial Economics*, 5: 147-175.

Ohlson, J. A., 1995, Earnings, Book Values, and Dividends in Equity Valuation, *Contemporary Accounting Research*, 661-687.

Palepu, Krishna G., Victor L. Bernard, and Paul M. Healy, 1996, Business Analysis & Valuation Using Financial Statements, South-Western College Publishing (Cincinnati, Ohio).

Stewart, G. Bennett, III, 1991, The Quest for Value: The EVA Management Guide, (Harper Business, New York, NY).

Exhibit 1: Numerical example of the relation between the value drivers and value creation

Combining equations 10 and 11 yields the following equation:

$$\text{Market/Book} = [(1+\gamma ROE)/(1+K_E)]^n + [ROE(1-\gamma)/(K_E - \gamma ROE)]\,[1 - \{(1+\gamma ROE)/(1+K_E)\}^n]$$

This Exhibit shows the hypothetical market-to-book ratios as a function of the three value drivers: profitability, advantage horizon, and re-investment.; assuming the firm has a cost of equity equal to 15%. The three cases differ by the level of reinvestment which varies from 0% to 66%.

Case #1: Reinvestment rate (γ) = 0%

	Return on Equity (ROE)		
Advantage Horizon	**5%**	**15%**	**25%**
5 years	0.66	1.00	1.34
15 years	0.42	1.00	1.58
30 years	0.34	1.00	1.66

Case #2: Reinvestment rate (γ) = 33%

	Return on Equity (ROE)		
Advantage Horizon	**5%**	**15%**	**25%**
5 years	0.65	1.00	1.39
15 years	0.37	1.00	1.88
30 years	0.27	1.00	2.24

Case #3: Reinvestment rate (γ) = 66%

	Return on Equity (ROE)		
Advantage Horizon	**5%**	**15%**	**25%**
5 years	0.65	1.00	1.45
15 years	0.32	1.00	2.43
30 years	0.18	1.00	4.27

Exhibit 2: Relation between Return on Equity (ROE) and Market-to-Book Ratio

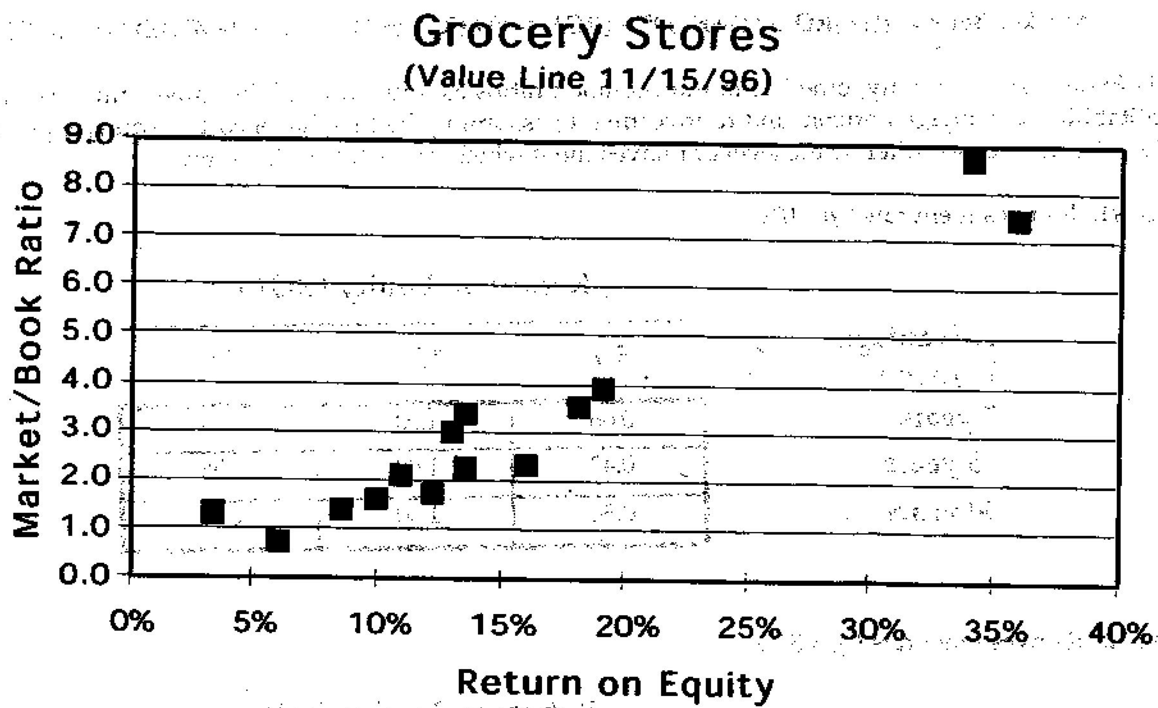

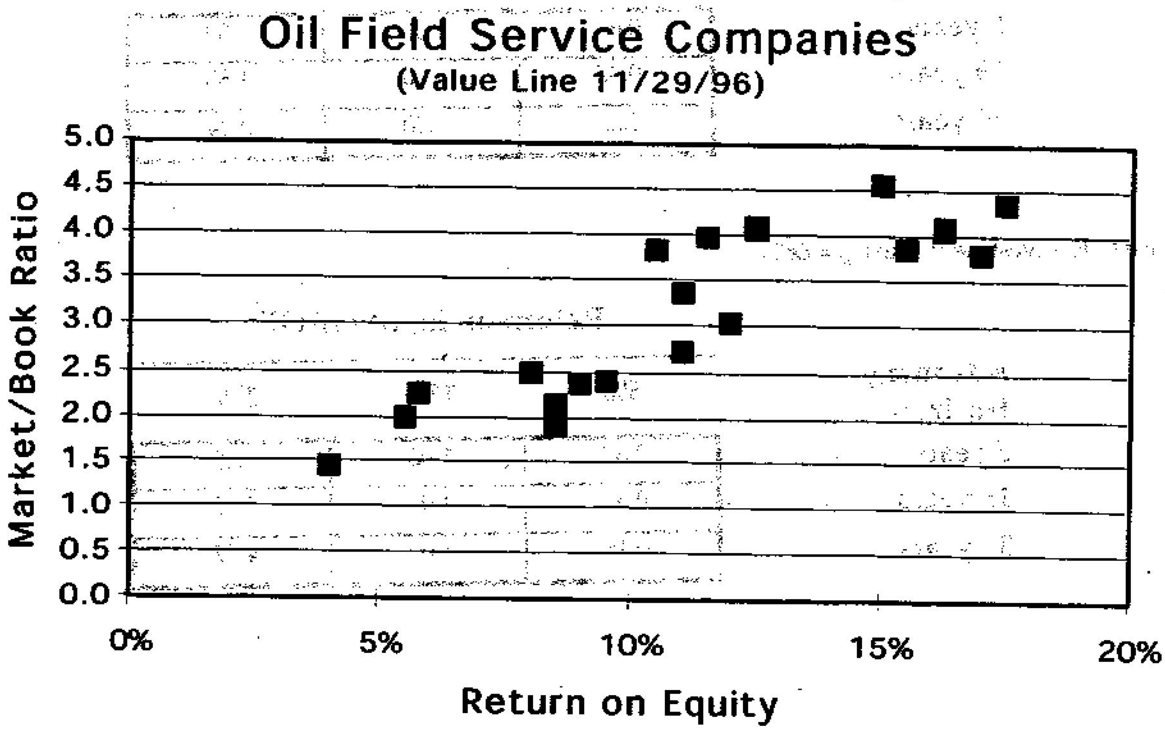

Exhibit 3: Advantage horizon

Fruhan (1995) analyzed the advantage horizon of a sample of 87 high-performing firms. To be included in the sample, firms had to have an average ROE of more than 25% for five consecutive years between 1976-82 and have sales greater than $200 million. He found the following:

Top Performers:	1976-78 Average ROE
1. Petrie Stores	2.03
2. H&R Block	1.45
3. Standard Microsystems	1.43
4. Airborne Freight	0.77
5. Wendy's International	0.74
6. Commerce Clearing House	0.67
7. Avon Products	0.67
8. Southwest Airlines	0.63
9. Charming Shoppes	0.56
10. Loctite Corp.	0.56

For the period from 1976-78:

Median ROE for the top 87 firms	= 37%
S&P 400 Average ROE	= 15%
Spread	= 21%

For the period from 1989-93:

Median ROE for the top 87 firms	= 17%
S&P 400 Average ROE	= 15%
Spread	= 2%

Lesson: The advantage horizon is finite.

Appendix 1: Equity value and the advantage horizon

Equations 1 and 3 show that a firm's equity market value is a function of its return on equity (ROE) and cost of equity (K_E). Assuming no retention of earnings and constant returns, equity value is:

$$E_{MV} = ROE^*E_{BV}\ /(1+K_E) + ROE^*E_{BV}\ /(1+K_E)^2 + \ldots \qquad \text{(A1.1)}$$

dividing through by the book value of equity (E_{BV}) yields

$$\text{Market/Book} = E_{MV}/E_{BV} = ROE/\ (1+K_E) + ROE\ /(1+K_E)^2 + \ldots \qquad \text{(A1.2)}$$

The ROE can be divided into two parts: ROE = $(ROE - K_E) + K_E$. The first term (ROE- KE) consists of "abnormal" earnings, returns to equity in excess of the cost of equity; the second term consists of "normal" earnings because that is the expected return on equity. Substituting back into equation A1.2 yields:

$$\text{Market/Book} = [ROE\text{-}K_E) + K_E]/(1+K_E) + [ROE\text{-}K_E) + K_E]\ /(1+K_E)^2 + \ldots \qquad \text{(A1.3)}$$

$$\text{Market/Book} = (ROE\text{-}K_E)/(1+K_E) + ([ROE\text{-}K_E)/(1+K_E)^2 + \ldots$$

$$+ K_E/(1+K_E) + K_E/(1+K_E)^2 + \ldots \qquad \text{(A1.4)}$$

Equation A1.4 is the sum of two geometric series, one of normal earnings and one of abnormal earnings. The present value of the normal earnings (using a perpetuity formula) is one:

$$1 = K_E/K_E = K_E/(1+K_E) + K_E/(1+K_E)^2 + \ldots \qquad \text{(A1.5)}$$

The present value of the abnormal earnings depends on how long the firm expects to earn abnormal earnings. It can be thought of as an annuity: The firm receives a stream of abnormal earnings for a period of n years. The present value of an annuity can be written as:

$$\text{present value} = (ROE\text{-}K_E) * [(1/K_E) - (1/(K_E(1+K_E)^n)] \qquad \text{(A1.6)}$$

Combining equations A1.5 and A1.6 yields:

$$\textbf{Market/Book} = 1 + \{(ROE\text{-}K_E)^*[(1/K_E) - (1/(K_E(1+K_E)^n)]\} \qquad \text{(A1.7)}$$

as n approaches infinity, equation A1.7 reduces to equation 4 in the note.

Appendix 2: Equity value and reinvestment

This appendix derives a model of equity valuation as a growing perpetuity. Given a firm with a constant return on equity (ROE), it can either retain its earnings or pay them out to equityholders as dividends. Assuming the firms retains a fraction of earnings (γ) and pays out the remainder, then the market value of equity can be determined as follows.

Time	Total Earnings	Amount Paid Out (ECF)	Amount Retained	Book Value of Equity
t=0				E_0
t=1	ROE^*E_0	$(1-\gamma)^*\ ROE^*E_0$	$(\gamma)^*\ ROE^*E_0$	$E_1 = E_0 + (\gamma)^*\ ROE^*E_0$ $E_1 = E_0\,(1+\gamma ROE)$
t=2	ROE^*E_1 $ROE^*[E_0(1+\gamma ROE)]$	$(1-\gamma)^*\ ROE^*E_1$ $(1-\gamma)^*\ ROE^*\ E_0\,(1+\gamma ROE)$	$(\gamma)^*\ ROE^*E_1$ $(\gamma)^*\ ROE^*\ E_0\,(1+\gamma ROE)$	$E_2 = E_1 + (\gamma)^*\ ROE^*E_1$ $E_2 = E_1\,(1+\gamma ROE)$ $E_2 = E_0\,(1+\gamma ROE)^2$
t=3	ROE^*E_2 $ROE^*[E_0(1+\gamma ROE)^2]$	$(1-\gamma)^*\ ROE^*E_2$ $(1-\gamma)^*\ ROE^*\ E_0\,(1+\gamma ROE)^2$	$(\gamma)^*\ ROE^*E_2$ $(\gamma)^*\ ROE^*\ E_0\,(1+\gamma ROE)^2$	$E_3 = E_2 + (\gamma)^*\ ROE^*E_2$ $E_3 = E_2\,(1+\gamma ROE)$ $E_3 = E_0\,(1+\gamma ROE)^3$
t=4	(etc.)			
Growth Rate	γROE	γROE	γROE	γROE

Value = discounted present value of payouts (equity cash flows)

$$= \frac{((1-\gamma)^*\ ROE^*E_0)}{(1+K_E)} + \frac{((1-\gamma)^*\ ROE^*\ E_0\,(1+\gamma ROE)}{(1+K_E)^2} + \ldots \qquad \text{(A2.1)}$$

$$= \frac{((1-\gamma)^*\ ROE^*E_0)}{(1+K_E)} \;\{1 + [(1+\gamma ROE)/(1+K_E)] + [(1+\gamma ROE)/(1+K_E)]^2 + \ldots\} \qquad \text{(A2.2)}$$

Equation A-2 is a growing perpetuity with growth rate equal to γROE. It can be rewritten as:

$$\text{Equity Value} = \frac{(1-\gamma)^*\ ROE^*E_0}{(K_E - \gamma ROE)} \qquad \text{(A2.3)}$$

After multiplying through by the book value of equity (E_0), one gets the ratio of equity at market value to equity at book value ($E_{MV}/E_{BV} = V/E_0$):

$$\textbf{Market/Book} = \frac{(1-\gamma)^*\ ROE}{(K_E - \gamma ROE)} \qquad \text{(A2.4)}$$

THE JOURNAL OF FINANCE • VOL. LVII, NO. 2 • APRIL 2002

The Equity Premium

EUGENE F. FAMA and KENNETH R. FRENCH*

ABSTRACT

We estimate the equity premium using dividend and earnings growth rates to measure the expected rate of capital gain. Our estimates for 1951 to 2000, 2.55 percent and 4.32 percent, are much lower than the equity premium produced by the average stock return, 7.43 percent. Our evidence suggests that the high average return for 1951 to 2000 is due to a decline in discount rates that produces a large unexpected capital gain. Our main conclusion is that the average stock return of the last half-century is a lot higher than expected.

THE EQUITY PREMIUM—the difference between the expected return on the market portfolio of common stocks and the risk-free interest rate—is important in portfolio allocation decisions, estimates of the cost of capital, the debate about the advantages of investing Social Security funds in stocks, and many other applications. The average return on a broad portfolio of stocks is typically used to estimate the expected market return. The average real return for 1872 to 2000 on the S&P index (a common proxy for the market portfolio, also used here) is 8.81 percent per year. The average real return on six-month commercial paper (a proxy for the risk-free interest rate) is 3.24 percent. This large spread (5.57 percent) between the average stock return and the interest rate is the source of the so-called equity premium puzzle: Stock returns seem too high given the observed volatility of consumption (Mehra and Prescott (1985)).

We use fundamentals (dividends and earnings) to estimate the expected stock return. Along with other evidence, the expected return estimates from fundamentals help us judge whether the realized average return is high or low relative to the expected value.

The logic of our approach is straightforward. The average stock return is the average dividend yield plus the average rate of capital gain:

$$A(R_t) = A(D_t/P_{t-1}) + A(GP_t), \tag{1}$$

* Fama is from the University of Chicago and French is from Dartmouth College. The comments of John Campbell, John Cochrane, Kent Daniel, John Heaton, Jay Ritter, Andrei Shleifer, Rex Sinquefield, Tuomo Vuolteenaho, Paul Zarowin, and seminar participants at Boston College, Dartmouth College, the NBER, Purdue University, the University of Chicago, and Washington University have been helpful. Richard Green (the editor) and the two referees get special thanks.

where D_t is the dividend for year t, P_{t-1} is the price at the end of year $t-1$, $GP_t = (P_t - P_{t-1})/P_{t-1}$ is the rate of capital gain, and $A(\ )$ indicates an average value. (Throughout the paper, we refer to D_t/P_{t-1} as the dividend yield and D_t/P_t is the dividend–price ratio. Similarly, Y_t/P_{t-1}, the ratio of earnings for year t to price at the end of year $t-1$, is the earnings yield and Y_t/P_t is the earnings–price ratio.)

Suppose the dividend–price ratio, D_t/P_t, is stationary (mean reverting). Stationarity implies that if the sample period is long, the compound rate of dividend growth approaches the compound rate of capital gain. Thus, an alternative estimate of the expected stock return is

$$A(RD_t) = A(D_t/P_{t-1}) + A(GD_t), \tag{2}$$

where $GD_t = (D_t - D_{t-1})/D_{t-1}$ is the growth rate of dividends. We call (2) the dividend growth model.

The logic that leads to (2) applies to any variable that is cointegrated with the stock price. For example, the dividend–price ratio may be nonstationary because firms move away from dividends toward share repurchases as a way of returning earnings to stockholders. But if the earnings–price ratio, Y_t/P_t, is stationary, the average growth rate of earnings, $A(GY_t) = A((Y_t - Y_{t-1})/Y_{t-1})$, is an alternative estimate of the expected rate of capital gain. And $A(GY_t)$ can be combined with the average dividend yield to produce another estimate of the expected stock return:

$$A(RY_t) = A(D_t/P_{t-1}) + A(GY_t). \tag{3}$$

We call (3) the earnings growth model.[1]

We should be clear about the expected return concept targeted by (1), (2), and (3). D_t/P_t and Y_t/P_t vary through time because of variation in the conditional (point-in-time) expected stock return and the conditional expected growth rates of dividends and earnings (see, e.g., Campbell and Shiller (1989)). But if the stock return and the growth rates are stationary (they have constant unconditional means), D_t/P_t and Y_t/P_t are stationary. Then, like the average return (1), the dividend and earnings growth models (2) and (3) provide estimates of the unconditional expected stock return. In short, the focus of the paper is estimates of the unconditional expected stock return.

The estimate of the expected real equity premium for 1872 to 2000 from the dividend growth model (2) is 3.54 percent per year. The estimate from the average stock return, 5.57 percent, is almost 60 percent higher. The difference between the two is largely due to the last 50 years. The equity premium for 1872 to 1950 from the dividend growth model, 4.17 percent per year, is close to the estimate from the average return, 4.40 percent. In con-

[1] Motivated by the model in Lettau and Ludvigson (2001), one can argue that if the ratio of consumption to stock market wealth is stationary, the average growth rate of consumption is another estimate of the expected rate of capital gain. We leave this path to future work.

trast, the equity premium for 1951 to 2000 produced by the average return, 7.43 percent per year, is almost three times the estimate, 2.55 percent, from (2). The estimate of the expected real equity premium for 1951 to 2000 from the earnings growth model (3), 4.32 percent per year, is larger than the estimate from the dividend growth model (2). But the earnings growth estimate is still less than 60 percent of the estimate from the average return.

Three types of evidence suggest that the lower equity premium estimates for 1951 to 2000 from fundamentals are closer to the expected premium. (a) The estimates from fundamentals are more precise. For example, the standard error of the estimate from the dividend growth model is less than half the standard error of the estimate from the average return. (b) The Sharpe ratio for the equity premium from the average stock return for 1951 to 2000 is just about double that for 1872 to 1950. In contrast, the equity premium from the dividend growth model has a similar Sharpe ratio for 1872 to 1950 and 1951 to 2000. (c) Most important, valuation theory specifies relations among the book-to-market ratio, the return on investment, and the cost of equity capital (the expected stock return). The estimates of the expected stock return for 1951 to 2000 from the dividend and earnings growth models line up with other fundamentals in the way valuation theory predicts. But the book-to-market ratio and the return on investment suggest that the expected return estimate from the average stock return is too high.

Our motivation for the dividend growth model (2) is simpler and more general, but (2) can be viewed as the expected stock return estimate of the Gordon (1962) model. Our work is thus in the spirit of a growing literature that uses valuation models to estimate expected returns (e.g., Blanchard (1993), Claus and Thomas (2001), and Gebhardt, Lee, and Swaminathan (2001)). Claus and Thomas and Gebhardt, Lee, and Swaminathan use forecasts by security analysts to estimate expected cash flows. Their analyst forecasts cover short periods (1985 to 1998 and 1979 to 1995). We use realized dividends and earnings from 1872 to 2000. This 129-year period provides a long perspective, which is important for judging the competing expected return estimates from fundamentals and realized stock returns. Moreover, though the issue is controversial (Keane and Runkle (1998)), Claus and Thomas find that analyst forecasts are biased; they tend to be substantially above observed growth rates. The average growth rates of dividends and earnings we use are unbiased estimates of expected growth rates.

Like us, Blanchard (1993) uses dividend growth rates to estimate the expected rate of capital gain, which he combines with an expected dividend yield to estimate the expected stock return. But his focus is different and his approach is more complicated than ours. He is interested in the path of the conditional expected stock return. His conditional expected return is the sum of the fitted values from time-series regressions of the realized dividend yield and a weighted average of 20 years of future dividend growth rates on four predetermined variables (the dividend yield, the real rate of capital gain, and the levels of interest rates and inflation). He focuses on describing the path of the conditional expected return in terms of his four explanatory variables.

In contrast, our prime interest is the unconditional expected return, which we estimate more simply as the sum of the average dividend yield and the average growth rate of dividends or earnings. This approach is valid if the dividend–price and earnings–price ratios are stationary. And we argue below that it continues to produce estimates of the average expected stock return when the price ratios are subject to reasonable forms of nonstationarity. Given its simplicity and generality, our approach is an attractive addition to the research toolbox for estimating the expected stock return.

Moreover, our focus is comparing alternative estimates of the unconditional expected stock return over the long 1872 to 2000 period, and explaining why the expected return estimates for 1951 to 2000 from fundamentals are much lower than the average return. Our evidence suggests that much of the high return for 1951 to 2000 is unexpected capital gain, the result of a decline in discount rates.

Specifically, the dividend–price and earnings–price ratios fall from 1950 to 2000; the cumulative percent capital gain for the period is more than three times the percent growth in dividends or earnings. All valuation models agree that the two price ratios are driven by expectations about future returns (discount rates) and expectations about dividend and earnings growth. Confirming Campbell (1991), Cochrane (1994), and Campbell and Shiller (1998), we find that dividend and earnings growth rates for 1950 to 2000 are largely unpredictable. Like Campbell and Shiller (1998), we thus infer that the decline in the price ratios is mostly due to a decline in expected returns. Some of this decline is probably expected, the result of reversion of a high 1950 conditional expected return to the unconditional mean. But most of the decline in the price ratios seems to be due to the unexpected decline of expected returns to ending values far below the mean.

The paper proceeds as follows. The main task, addressed in Sections I and II, is to compare and evaluate the estimates of the unconditional annual expected stock return provided by the average stock return and the dividend and earnings growth models. Section III then considers the issues that arise if the goal is to estimate the long-term expected growth of wealth, rather than the unconditional expected annual (simple) return. Section IV concludes.

I. The Unconditional Annual Expected Stock Return

Table I shows estimates of the annual expected real equity premium for 1872 to 2000. The market portfolio is the S&P 500 and its antecedents. The deflator is the Producer Price Index until 1925 (from Shiller (1989)) and the Consumer Price Index thereafter (from Ibbotson Associates). The risk-free interest rate is the annual real return on six-month commercial paper, rolled over at midyear. The risk-free rate and S&P earnings data are from Shiller, updated by Vuolteenaho (2000) and us. Beginning in 1925, we construct S&P book equity data from the book equity data in Davis, Fama, and French (2000), expanded to include all NYSE firms. The data on dividends, prices, and returns for 1872 to 1925 are from Shiller. Shiller's annual data on the

Table I

Real Equity Premium and Related Statistics for the S&P Portfolio

The inflation rate for year t is $Inf_t = L_t/L_{t-1} - 1$, where L_t is the price level at the end of year t. The real return for year t on six-month (three-month for the year 2000) commercial paper (rolled over at midyear) is F_t. The nominal values of book equity and price for the S&P index at the end of year t are b_t and p_t. Nominal S&P dividends and earnings for year t are d_t and y_t. Real rates of growth of dividends, earnings, and the stock price are $GD_t = (d_t/d_{t-1})*(L_{t-1}/L_t) - 1$, $GY_t = (y_t/y_{t-1})*(L_{t-1}/L_t) - 1$, and $GP_t = (p_t/p_{t-1})*(L_{t-1}/L_t) - 1$. The real dividend yield is $D_t/P_{t-1} = (d_t/p_{t-1})*(L_{t-1}/L_t)$. The real income return on investment is $Y_t/B_{t-1} = (1 + y_t/b_{t-1})*(L_{t-1}/L_t) - 1$. The dividend growth estimate of the real S&P return for t is $RD_t = D_t/P_{t-1} + GD_t$, the earnings growth estimate is $RY_t = D_t/P_{t-1} + GY_t$, and R_t is the realized real S&P return. The dividend and earnings growth estimates of the real equity premium for year t are $RXD_t = RD_t - F_t$ and $RXY_t = RY_t - F_t$, and $RX_t = R_t - F_t$ is the real equity premium from the realized real return. The Sharpe ratio for $RD_t - F_t$ (the mean of $RD_t - F_t$ divided by the standard deviation of R_t) is SD, SY is the Sharpe ratio for $RY_t - F_t$ (the mean of $RY_t - F_t$ divided by the standard deviation of R_t), and SR is the Sharpe ratio for $R_t - F_t$ (the mean of $R_t - F_t$ divided by the standard deviation of R_t). Except for the Sharpe ratios, all variables are expressed as percents, that is, they are multiplied by 100.

	Inf_t	F_t	D_t/P_{t-1}	GD_t	GY_t	GP_t	RD_t	RY_t	R_t	RXD_t	RXY_t	RX_t	SD	SY	SR
Means of annual values of variables															
1872–2000	2.16	3.24	4.70	2.08	NA	4.11	6.78	NA	8.81	3.54	NA	5.57	0.20	NA	0.31
1872–1950	0.99	3.90	5.34	2.74	NA	2.96	8.07	NA	8.30	4.17	NA	4.40	0.22	NA	0.23
1951–2000	4.00	2.19	3.70	1.05	2.82	5.92	4.74	6.51	9.62	2.55	4.32	7.43	0.15	0.25	0.44
Standard deviations of annual values of variables															
1872–2000	7.51	8.48	1.39	12.37	NA	17.83	12.56	NA	18.03	13.00	NA	18.51			
1872–1950	9.11	10.63	1.12	15.28	NA	18.48	15.41	NA	18.72	16.02	NA	19.57			
1951–2000	3.11	2.46	1.17	5.09	13.79	16.77	5.21	13.51	17.03	5.62	14.02	16.73			
Means of annual continuously compounded returns and growth rates															
1872–2000	1.86	2.87		1.34	NA	2.48			7.00						
1872–1950	0.59	3.33		1.60	NA	1.22			6.41						
1951–2000	3.88	2.14		0.92	1.89	4.46			7.94						

	b_t/p_t	RD_t	RY_t	R_t	Y_t/B_{t-1}
Means of annual values of variables					
1951–2000	0.66	4.74	6.51	9.62	7.60

level of the S&P (used to compute returns and other variables involving price) are averages of daily January values. The S&P dividend, price, and return data for 1926 to 2000 are from Ibbotson Associates, and the returns for 1926 to 2000 are true annual returns.

Without showing the details, we can report that the CRSP value-weight portfolio of NYSE, AMEX, and Nasdaq stocks produces average returns and dividend growth estimates of the expected return close to the S&P estimates for periods after 1925 when both indices are available. What one takes to be the risk-free rate has a bigger effect. For example, substituting the one-month Treasury bill rate for the six-month commercial paper rate causes estimates of the annual equity premium for 1951 to 2000 to rise by about one percent. But for our main task—comparing equity premium estimates from (1), (2), and (3)—differences in the risk-free rate are an additive constant that does not affect inferences.

One can estimate expected returns in real or nominal terms. Since portfolio theory says the goal of investment is consumption, real returns seem more relevant, and only results for real returns are shown. Because of suspicions about the quality of the price deflator during the early years of 1872 to 2000, we have replicated the results for nominal returns. They support all the inferences from real returns.

The dividend and earnings growth models (2) and (3) assume that the market dividend–price and earnings–price ratios are stationary. The first three annual autocorrelations of D_t/P_t for 1872 to 2000 are 0.73, 0.51, and 0.47. For the 1951 to 2000 period that occupies much of our attention, the autocorrelations are 0.83, 0.72, and 0.69. The autocorrelations are large, but their decay is roughly like that of a stationary first-order autoregression (AR1). This is in line with formal evidence (Fama and French (1988), Cochrane (1994), and Lamont (1998)) that the market dividend–price ratio is highly autocorrelated but slowly mean-reverting. S&P earnings data for the early years of 1872 to 2000 are of dubious quality (Shiller (1989)), so we estimate expected returns with the earnings growth model (3) only for 1951 to 2000. The first three autocorrelations of Y_t/P_t for 1951 to 2000, 0.80, 0.70, and 0.61, are again roughly like those of a stationary AR1.

We emphasize, however, that our tests are robust to reasonable nonstationarity of D_t/P_t and Y_t/P_t. It is not reasonable that the expected stock return and the expected growth rates of dividends and earnings that drive D_t/P_t and Y_t/P_t are nonstationary processes that can wander off to infinity. But nonstationarity of D_t/P_t and Y_t/P_t due to structural shifts in productivity or preferences that permanently change the expected return or the expected growth rates is reasonable. Such regime shifts are not a problem for the expected return estimates from (2) and (3), as long as D_t/P_t and Y_t/P_t mean-revert within regimes. If the regime shift is limited to expected dividend and earnings growth rates, the permanent change in expected growth rates is offset by a permanent change in the expected dividend yield, and (2) and (3) continue to estimate the (stationary) expected stock return. (An Appendix, available on request, provides an example.) If there is a permanent shift in the expected stock return, it is nonstationary, but like the average return in (1), the dividend and earnings growth models in (2) and (3) estimate the average expected return during the sample period.

Indeed, an advantage of the expected return estimates from fundamentals is that they are likely to be less sensitive than the average return to long-lived shocks to dividend and earnings growth rates or the expected stock return. For example, a permanent shift in the expected return affects the average dividend yield, which is common to the three expected return estimates, but it produces a shock to the capital gain term in the average return in (1) that is not shared by the estimates in (2) and (3). In short, the estimates of the expected stock return from fundamentals are likely to be more precise than the average stock return.

A. The Equity Premium

For much of the period from 1872 to 2000—up to about 1950—the dividend growth model and the average stock return produce similar estimates of the expected return. Thereafter, the two estimates diverge. To illustrate, Table I shows results for 1872 to 1950 (79 years) and 1951 to 2000 (50 years). The year 1950 is a big year, with a high real stock return (23.40 percent), and high dividend and earnings growth estimates of the return (29.96 percent and 24.00 percent). But because the three estimates of the 1950 return are similarly high, the ordering of expected return estimates, and the inferences we draw from them, are unaffected by whether 1950 is allocated to the earlier or the later period. Indeed, pushing the 1950 break-year backward or forward several years does not affect our inferences.

For the earlier 1872 to 1950 period, there is not much reason to favor the dividend growth estimate of the expected stock return over the average return. Precision is not an issue; the standard errors of the two estimates are similar (1.74 percent and 2.12 percent), the result of similar standard deviations of the annual dividend growth rate and the rate of capital gain, 15.28 percent and 18.48 percent. Moreover, the dividend growth model and the average return provide similar estimates of the expected annual real return for 1872 to 1950, 8.07 percent and 8.30 percent. Given similar estimates of the expected return, the two approaches produce similar real equity premiums for 1872 to 1950, 4.17 percent (dividend growth model) and 4.40 percent (stock returns).

The competition between the dividend growth model and the average stock return is more interesting for 1951 to 2000. The dividend growth estimate of the 1951 to 2000 expected return, 4.74 percent, is less than half the average return, 9.62 percent. The dividend growth estimate of the equity premium, 2.55 percent, is 34 percent of the estimate from returns, 7.43 percent. The 1951 to 2000 estimates of the expected stock return and the equity premium from the earnings growth model, 6.51 percent and 4.32 percent, are higher than for the dividend growth model. But they are well below the estimates from the average return, 9.62 percent and 7.43 percent.

B. Evaluating the Expected Return Estimates for 1951 to 2000

We judge that the estimates of the expected stock return for 1951 to 2000 from fundamentals are closer to the true expected value, for three reasons.

(a) The expected return estimates from the dividend and earnings growth models are more precise than the average return. The standard error of the dividend growth estimate of the expected return for 1951 to 2000 is 0.74 percent, versus 2.43 percent for the average stock return. Since earnings growth is more volatile than dividend growth, the standard error of the expected return from the earnings growth model, 1.93 percent, is higher than the estimate from the dividend growth model, but it is smaller than the 2.43 percent standard error of the average stock return. Claus and Thomas (2001) also argue that expected return estimates from fundamentals are more precise than average returns, but they provide no direct evidence.

(b) Table I shows Sharpe ratios for the three equity premium estimates. Only the average premium in the numerator of the Sharpe ratio differs for the three estimates. The denominator for all three is the standard deviation of the annual stock return. The Sharpe ratio for the dividend growth estimate of the equity premium for 1872 to 1950, 0.22, is close to that produced by the average stock return, 0.23. More interesting, the Sharpe ratio for the equity premium for 1951 to 2000 from the dividend growth model, 0.15, is lower than but similar to that for 1872 to 1950. The Sharpe ratio for the 1951 to 2000 equity premium from the earnings growth model, 0.25, is somewhat higher than the dividend growth estimate, 0.15, but it is similar to the estimates for 1872 to 1950 from the dividend growth model, 0.22, and the average return, 0.23.

In asset pricing theory, the Sharpe ratio is related to aggregate risk aversion. The Sharpe ratios for the 1872 to 1950 and 1951 to 2000 equity premiums from the dividend growth model and the earnings growth model suggest that aggregate risk aversion is roughly similar in the two periods. In contrast, though return volatility falls a bit, the equity premium estimate from the average stock return increases from 4.40 percent for 1872 to 1950 to 7.43 percent for 1951 to 2000, and its Sharpe ratio about doubles, from 0.23 to 0.44. It seems implausible that risk aversion increases so much from the earlier to the later period.

(c) Most important, the behavior of other fundamentals favors the dividend and earnings growth models. The average ratio of the book value of equity to the market value of equity for 1951 to 2000 is 0.66, the book-to-market ratio B_t/P_t is never greater than 1.12, and it is greater than 1.0 for only 6 years of the 50-year period. Since, on average, the market value of equity is substantially higher than its book value, it seems safe to conclude that, on average, the expected return on investment exceeds the cost of capital.

Suppose investment at time $t-1$ generates a stream of equity earnings for $t, t+1, \ldots, t+N$ with a constant expected value. The average income return on book equity, $A(Y_t/B_{t-1})$, is then an estimate of the expected return on equity's share of assets. It is an unbiased estimate when N is infinite and it is upward biased when N is finite. In either case, if the expected return on investment exceeds the cost of capital, we should find that (except for sampling error) the average income return on book equity is greater than estimates of the cost of equity capital (the expected stock return):

$$A(Y_t/B_{t-1}) > E(R). \tag{4}$$

Table I shows that (4) is confirmed when we use the dividend and earnings growth models to estimate the expected real stock return for 1951 to 2000. The estimates of $E(R)$, 4.74 percent (dividend growth model) and 6.51 percent (earnings growth model), are below 7.60 percent, the average real income return on book equity, $A(Y_t/B_{t-1})$. In contrast, the average real stock return for 1951 to 2000, 9.62 percent, exceeds the average income return by more than 2 percent. An expected stock return that exceeds the expected income return on book equity implies that the typical corporate investment has a negative net present value. This is difficult to reconcile with an average book-to-market ratio substantially less than one.

To what extent are our results new? Using analyst forecasts of expected cash flows and a more complicated valuation model, Claus and Thomas (2001) produce estimates of the expected stock return for 1985 to 1998 far below the average return. Like us, they argue that the estimates from fundamentals are closer to the true expected return. We buttress this conclusion with new results on three fronts. (a) The long-term perspective provided by the evidence that, for much of the 1872 to 2000 period, average returns and fundamentals produce similar estimates of the expected return. (b) Direct evidence that the expected return estimates for 1951 to 2000 from fundamentals are more precise. (c) Sharpe ratios and evidence on how the alternative expected return estimates line up with the income return on investment. These new results provide support for the expected return estimates from fundamentals, and for the more specific inference that the average stock return for 1951 to 2000 is above the expected return.

II. Unexpected Capital Gains

Valuation theory suggests three potential explanations for why the 1951 to 2000 average stock return is larger than the expected return. (a) Dividend and earnings growth for 1951 to 2000 is unexpectedly high. (b) The expected (post-2000) growth rates of dividends and earnings are unexpectedly high. (c) The expected stock return (the equity discount rate) is unexpectedly low at the end of the sample period.

A. Is Dividend Growth for 1951 to 2000 Unexpectedly High?

If the prosperity of the United States over the last 50 years was not fully anticipated, dividend and earnings growth for 1951 to 2000 exceed 1950 expectations. Such unexpected in-sample growth produces unexpected capital gains. But it does not explain why the average return for 1951 to 2000 (the average dividend yield plus the average rate of capital gain) is so much higher than the expected return estimates from fundamentals (the average dividend yield plus the average growth rate of dividends or earnings). To see the point, note that unexpected in-sample dividend and earnings growth do not affect either the 1950 or the 2000 dividend–price and earnings–price ratios. (The 2000 ratios depend on post-2000 expected returns and growth rates.) Suppose D_t/P_t and E_t/P_t were the same in 1950 and 2000. Then the total percent growth in dividends and earnings during the period would be the same as the percent growth in the stock price. And (1), (2), and (3) would provide similar estimates of the expected stock return.

It is worth dwelling on this point. There is probably survivor bias in the U.S. average stock return for 1872 to 1950, as well as for 1951 to 2000. During the 1872 to 2000 period, it was not a foregone conclusion that the U.S. equity market would survive several financial panics, the Great Depression, two world wars, and the cold war. The average return for a market that survives many potentially cataclysmic challenges is likely to be higher than the expected return (Brown, Goetzmann, and Ross (1995)). But if the positive bias shows up only as higher than expected dividend and earnings growth during the sample period, there is similar survivor bias in the expected return estimates from fundamentals—a problem we do not solve. Our more limited goal is to explain why the average stock return for 1951 to 2000 is so high relative to the expected return estimates from the dividend and earnings growth models.

Since unexpected growth for 1951 to 2000 has a similar effect on the three expected return estimates, the task of explaining why the estimates are so different falls to the end-of-sample values of future expected returns and expected dividend and earnings growth. We approach the problem by first looking for evidence that expected dividend or earnings growth is high at the end of the sample period. We find none. We then argue that the large spread of capital gains over dividend and earnings growth for 1951 to 2000, or equivalently, the low end-of-sample dividend–price and earnings–price ratios, are due to an unexpected decline in expected stock returns to unusually low end-of-sample values.

B. Are Post-2000 Expected Dividend and Earnings Growth Rates Unusually High?

The behavior of dividends and earnings provides little evidence that rationally assessed (i.e., true) long-term expected growth is high at the end of the sample period. If anything, the growth rate of real dividends declines during the 1951 to 2000 period (Table II). The average growth rate for the first two decades, 1.60 percent, is higher than the average growth rates for the last three, 0.68 percent. The regressions in Table III are more formal evidence on the best forecast of post-2000 real dividend growth rates. Re-

Table II

Means of Simple Real Equity Premium and Related Statistics for the S&P Portfolio for 10-year Periods

The inflation rate for year t is $Inf_t = L_t/L_{t-1} - 1$, where L_t is the price level at the end of year t. The real return for year t on six-month (three-month for the year 2000) commercial paper (rolled over at midyear) is F_t. The nominal price of the S&P index at the end of year t is p_t. Nominal S&P dividends and earnings for year t are d_t and y_t. Real rates of growth of dividends, earnings, and the stock price are $GD_t = (d_t/d_{t-1})*(L_{t-1}/L_t) - 1$, $GY_t = (y_t/y_{t-1})*(L_{t-1}/L_t) - 1$, and $GP_t = (p_t/p_{t-1})*(L_{t-1}/L_t) - 1$. The real dividend yield is $D_t/P_{t-1} = (d_t/p_{t-1})*(L_{t-1}/L_t)$. The dividend growth estimate of the real S&P return for t is $RD_t = D_t/P_{t-1} + GD_t$, the earnings growth estimate is $RY_t = D_t/P_{t-1} + GY_t$, and R_t is the realized real S&P return. The dividend and earnings growth estimates of the real equity premium for year t are $RXD_t = RD_t - F_t$ and $RXY_t = RY_t - F_t$, and $RX_t = R_t - F_t$ is the real equity premium from the realized real return. All variables are expressed as percents, that is, they are multiplied by 100.

	Inf_t	F_t	D_t/P_{t-1}	GD_t	GY_t	GP_t	RD_t	RY_t	R_t	RXD_t	RXY_t	RX_t
1872–1880	−2.77	9.86	6.29	4.62	NA	7.13	10.91	NA	13.42	1.06	NA	3.56
1881–1890	−1.72	7.23	5.04	0.69	NA	0.04	5.73	NA	5.08	−1.51	NA	−2.15
1891–1900	0.18	5.08	4.40	4.49	NA	4.75	8.89	NA	9.15	3.81	NA	4.08
1901–1910	1.95	3.18	4.45	3.25	NA	2.33	7.70	NA	6.78	4.52	NA	3.60
1911–1920	6.82	0.82	5.70	−3.43	NA	−6.52	2.27	NA	−0.83	1.45	NA	−1.64
1921–1930	−1.70	7.41	5.72	9.07	NA	11.83	14.78	NA	17.54	7.37	NA	10.13
1931–1940	−1.23	2.80	5.31	0.36	NA	2.21	5.67	NA	7.52	2.87	NA	4.72
1941–1950	6.04	−4.57	5.90	3.02	NA	2.33	8.91	NA	8.22	13.48	NA	12.79
1951–1960	1.79	1.05	4.68	1.22	0.61	10.64	5.90	5.30	15.32	4.85	4.24	14.27
1961–1970	2.94	2.27	3.21	1.98	2.07	2.69	5.19	5.27	5.90	2.92	3.01	3.63
1971–1980	8.11	−0.30	4.04	−0.86	3.47	−1.92	3.18	7.50	2.12	3.48	7.80	2.42
1981–1990	4.51	5.32	4.19	2.32	0.37	5.40	6.51	4.56	9.59	1.19	−0.75	4.28
1991–2000	2.68	2.61	2.36	0.58	7.58	12.80	2.94	9.94	15.16	0.32	7.32	12.54

Table III

Regressions to Forecast Real Dividend and Earnings Growth Rates, GD_t and GY_t

The price level at the end of year t is L_t. The nominal values of book equity and price for the S&P index at the end of year t are b_t and p_t. Nominal S&P dividends and earnings for year t are d_t and y_t. The real dividend and earnings growth rates for year t are $GD_t = (d_t/d_{t-1})*(L_{t-1}/L_t) - 1$ and $GY_t = (y_t/y_{t-1})*(L_{t-1}/L_t) - 1$, and R_t is the realized real return on the S&P portfolio for year t. The regression intercept is *Int*, and t-Stat is the regression coefficient (*Coef*) divided by its standard error. The regression R^2 is adjusted for degrees of freedom. Except for the dividend payout ratio, d_t/y_t, all variables are expressed as percents, that is, they are multiplied by 100.

Panel A: One Year: The Regressions Forecast Real Dividend Growth, GD_t, with Variables Known at $t - 1$

	Int	d_{t-1}/y_{t-1}	d_{t-1}/p_{t-1}	GD_{t-1}	GD_{t-2}	GD_{t-3}	R_{t-1}	R_{t-2}	R_{t-3}	R^2
1875–1950, N = 76 years										
Coef	29.56	−23.12	−2.63	−0.12	−0.07	−0.03	0.22	0.13	0.09	0.38
t-Stat	3.22	−3.17	−1.77	−1.08	−0.64	−0.29	2.24	1.37	1.01	
1951–2000, N = 50 years										
Coef	−2.16	2.97	0.11	−0.07	−0.20	−0.06	0.11	0.07	0.01	0.01
t-Stat	−0.40	0.33	0.16	−0.45	−1.57	−0.45	2.17	1.33	0.22	

Panel B: Two Years: The Regressions Forecast Real Dividend Growth, GD_t, with Variables Known at $t-2$

	Int	d_{t-2}/y_{t-2}	d_{t-2}/p_{t-2}	GD_{t-2}	GD_{t-3}	R_{t-2}	R_{t-3}	R^2
1875–1950, N = 76 years								
Coef	6.61	−11.60	0.31	−0.26	0.05	0.24	0.11	0.07
t-Stat	0.64	−1.28	0.18	−2.02	0.39	2.03	1.00	
1951–2000, N = 50 years								
Coef	−4.11	7.62	0.32	−0.14	−0.03	0.05	−0.01	−0.05
t-Stat	−0.73	0.81	0.46	−1.13	−0.28	0.99	−0.16	

Panel C: One Year: The Regressions Forecast Real Earnings Growth, GY_t, with Variables Known at $t-1$

	Int	Y_{t-1}/B_{t-2}	d_{t-1}/y_{t-1}	y_{t-1}/p_{t-1}	GY_{t-1}	GY_{t-2}	GY_{t-3}	R_{t-1}	R_{t-2}	R_{t-3}	R^2
1951–2000, N = 50 years											
Coef	5.48	0.11	13.06	−1.36	0.21	−0.13	−0.31	0.28	−0.25	0.03	0.40
t-Stat	0.33	0.11	0.52	−1.91	1.17	−0.89	−2.64	2.39	−2.18	0.26	

Panel D: Two Years: The Regressions Forecast Real Earnings Growth, GY_t, with Variables Known at $t-2$

	Int	Y_{t-2}/B_{t-3}	d_{t-2}/y_{t-2}	y_{t-2}/p_{t-2}	GY_{t-2}	GY_{t-3}	R_{t-2}	R_{t-3}	R^2
1951–2000, N = 50 years									
Coef	−7.60	0.46	2.05	−0.74	−0.16	−0.39	−0.31	−0.12	0.23
t-Stat	−0.43	1.66	0.76	−1.02	−0.92	−2.54	−2.59	−0.97	

gressions are shown for forecasts one year ahead (the explanatory variables for year t dividend growth are known at the end of year $t - 1$) and two years ahead (the explanatory variables are known at the end of year $t - 2$).

The regression for 1875 to 1950 suggests strong forecast power one year ahead. The slopes on the lagged payout ratio, the dividend–price ratio, and the stock return are close to or more than two standard errors from zero, and the regression captures 38 percent of the variance of dividend growth. Even in the 1875 to 1950 period, however, power to forecast dividend growth does not extend much beyond a year. When dividend growth for year t is explained with variables known at the end of year $t - 2$, the regression R^2 falls from 0.38 to 0.07. Without showing the details, we can report that extending the forecast horizon from two to three years causes all hint of forecast power to disappear. Thus, for 1875 to 1950, the best forecast of dividend growth more than a year or two ahead is the historical average growth rate.

We are interested in post-2000 expected dividend growth, and even the short-term forecast power of the dividend regressions for 1872 to 1950 evaporates in the 1951 to 2000 period. The lagged stock return has some information ($t = 2.17$) about dividend growth one year ahead. But the 1951 to 2000 regression picks up only one percent of the variance of dividend growth. And forecast power does not improve for longer forecast horizons. Our evidence that dividend growth is essentially unpredictable during the last 50 years confirms the results in Campbell (1991), Cochrane (1991, 1994), and Campbell and Shiller (1998). If dividend growth is unpredictable, the historical average growth rate is the best forecast of future growth.

Long-term expected earnings growth also is not unusually high in 2000. There is no clear trend in real earnings growth during the 1951 to 2000 period. The most recent decade, 1991 to 2000, produces the highest average growth rate, 7.58 percent per year (Table II). But earnings growth is volatile. The standard errors of 10-year average growth rates vary around 5 percent. It is thus not surprising that 1981 to 1990, the decade immediately preceding 1991 to 2000, produces the lowest average real earnings growth rate, 0.37 percent per year.

The regressions in Table III are formal evidence on the predictability of earnings growth during the 1951 to 2000 period. There is some predictability of near-term growth, but it is largely due to transitory variation in earnings that is irrelevant for forecasting long-term earnings. In the 1951 to 2000 regression to forecast earnings growth one year ahead, the slope on the first lag of the stock return is positive (0.28, $t = 2.39$), but the slope on the second lag is negative (-0.25, $t = -2.18$) and about the same magnitude. Thus, the prediction of next year's earnings growth from this year's return is reversed the following year. In the one-year forecast regression for 1951 to 2000, the only variable other than lagged returns with power to forecast earnings growth ($t = -2.64$) is the third lag of earnings growth. But the slope is negative, so it predicts that the strong earnings growth of recent years is soon to be reversed.

In the 1951 to 2000 regression to forecast earnings one year ahead, there is a hint ($t = -1.91$) that the low earnings–price ratio at the end of the period implies higher than average expected growth one year ahead. But the effect peters out quickly; the slope on the lagged earnings–price ratio in the regression to forecast earnings growth two years ahead is -1.02 standard errors from zero. The only variables with forecast power two years ahead are the second lag of the stock return and the third lag of earnings growth. But the slopes on these variables are negative, so again the 2000 prediction is that the strong earnings growth of recent years is soon to be reversed. And again, regressions (not shown) confirm that forecast power for 1951 to 2000 does not extend beyond two years. Thus, beyond two years, the best forecast of earnings growth is the historical average growth rate.

In sum, the behavior of dividends for 1951 to 2000 suggests that future growth is largely unpredictable, so the historical mean growth rate is a near optimal forecast of future growth. Earnings growth for 1951 to 2000 is somewhat predictable one and two years ahead, but the end-of-sample message is that the recent high growth rates are likely to revert quickly to the historical mean. It is also worth noting that the market survivor bias argument of Brown, Goetzmann, and Ross (1995) suggests that past average growth rates are, if anything, upward biased estimates of future growth. In short, we find no evidence to support a forecast of strong future dividend or earnings growth at the end of our sample period.

C. Do Expected Stock Returns Fall during the 1951 to 2000 Period?

The S&P dividend–price ratio, D_t/P_t, falls from 7.18 percent at the end of 1950 to a historically low 1.22 percent at the end of 2000 (Figure 1). The growth in the stock price, P_{2000}/P_{1950}, is thus 5.89 times the growth in dividends, D_{2000}/D_{1950}. The S&P earnings–price ratio, Y_t/P_t, falls from 13.39 percent at the end of 1950 to 3.46 percent at the end of 2000, so the percent capital gain of the last 50 years is 3.87 times the percent growth in earnings. (Interestingly, almost all of the excess capital gain occurs in the last 20 years; Figure 1 shows that the 1979 earnings–price ratio, 13.40 percent, is nearly identical to the 13.39 percent value of 1950.)

All valuation models say that D_t/P_t and E_t/P_t are driven by expected future returns (discount rates) and expectations about future dividend and earnings growth. Our evidence suggests that rational forecasts of long-term dividend and earnings growth rates are not unusually high in 2000. We conclude that the large spread of capital gains for 1951 to 2000 over dividend and earnings growth is largely due to a decline in the expected stock return.

Some of the decline in D_t/P_t and E_t/P_t during 1951 to 2000 is probably anticipated in 1950. The dividend–price ratio for 1950, 7.18 percent, is high (Figure 1). The average for 1872 to 2000 is 4.64 percent. If D_t/P_t is mean-reverting, the expectation in 1950 of the yield in 2000 is close to the unconditional mean, say 4.64 percent. The actual dividend–price ratio for 2000 is

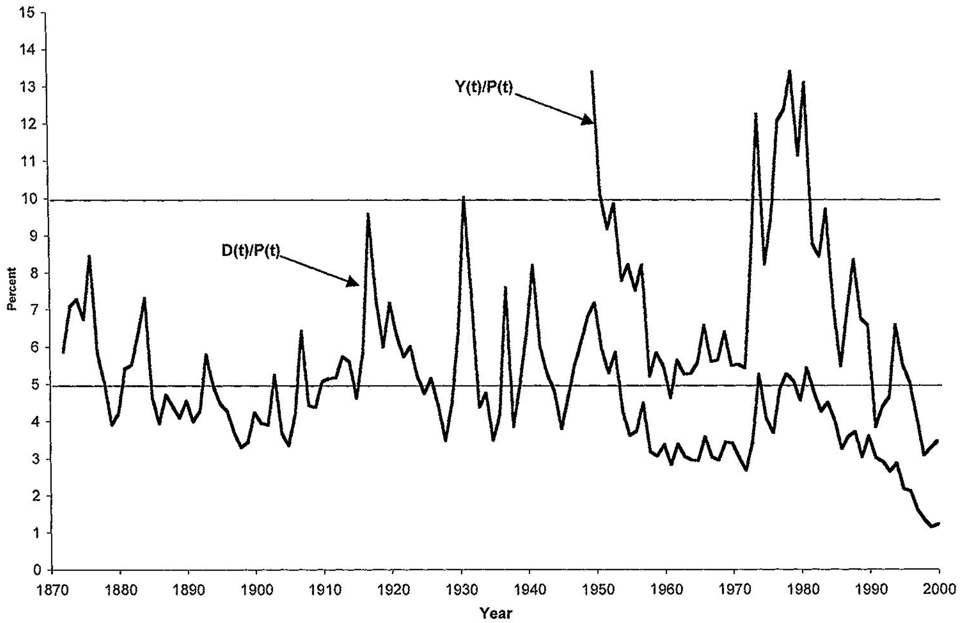

Figure 1. Dividend–price and earnings–price ratios.

1.22 percent. The 2000 stock price is thus 4.64/1.22 = 3.80 times what it would be if the dividend yield for 2000 hit the historical mean. Roughly speaking, this unexpected capital gain adds about 2.67 percent to the compound annual return for 1951 to 2000.

Similarly, part of the large difference between the 1951 to 2000 capital gain and the growth in earnings is probably anticipated in 1950. The 13.39 percent value of Y_t/P_t in 1950 is high relative to the mean for 1951 to 2000, 7.14 percent. If the earnings–price ratio is stationary, the expectation in 1950 of Y_t/P_t for 2000 is close to the unconditional mean, say 7.14 percent. The actual Y_t/P_t for 2000 is 3.46 percent. Thus, the 2000 stock price is 7.14/3.46 = 2.06 times what it would be if the ratio for 2000 hit the 7.14 percent average value for 1951 to 2000. Roughly speaking, this estimate of the unexpected capital gain adds about 1.45 percent to the compound annual return for the 50-year period.

In short, the percent capital gain for 1951 to 2000 is several times the growth of dividends or earnings. The result is historically low dividend–price and earnings–price ratios at the end of the period. Since the ratios are high in 1950, some of their subsequent decline is probably expected, but much of it is unexpected. Given the evidence that rational forecasts of long-term growth rates of dividends and earnings are not high in 2000, we conclude that the unexpected capital gains for 1951 to 2000 are largely due to a decline in the discount rate. In other words, the low end-of-sample price ratios imply low (rationally assessed, or true) expected future returns.

Like us, Campbell (1991), Cochrane (1994), and Campbell and Shiller (1998) find that, for recent periods, dividend and earnings growth are largely unpredictable, so variation in dividend–price and earnings–price ratios is largely due to the expected stock return. The samples in Campbell (1991) and Cochrane (1994) end in 1988 (before the strong subsequent returns that produce sharp declines in the price ratios), and they focus on explaining, in general terms, how variation in D_t/P_t splits between variation in the expected stock return and expected dividend growth. Campbell and Shiller (1998) focus on the low expected future returns implied by the low price ratios of recent years.

In contrast, we are more interested in what the decline in the price ratios says about past returns, specifically, that the average return for 1951 to 2000 is above the expected return. And this inference does not rest solely on the information in price ratios. We buttress it with two types of novel evidence. (a) The perspective from our long sample period that, although the average stock return for 1951 to 2000 is much higher than expected return estimates from fundamentals, the two approaches produce similar estimates for 1872 to 1950. (b) Evidence from Sharpe ratios, the book-to-market ratio, and the income return on investment, which also suggests that the average return for 1951 to 2000 is above the expected value.

III. Estimating the Expected Stock Return: Issues

There are two open questions about our estimates of the expected stock return. (a) In recent years the propensity of firms to pay dividends declines and stock repurchases surge. How do these changes in dividend policy affect our estimates of the expected return? (b) Under rather general conditions, the dividend and earnings growth models (2) and (3) provide estimates of the expected stock return. Are the estimates biased and does the bias depend on the return horizon? This section addresses these issues.

A. Repurchases and the Declining Incidence of Dividend Payers

Share repurchases surge after 1983 (Bagwell and Shoven (1989) and Dunsby (1995)), and, after 1978, the fraction of firms that do not pay dividends steadily increases (Fama and French (2001)). More generally, dividends are a policy variable, and changes in policy can raise problems for estimates of the expected stock return from the dividend growth model. There is no problem in the long-term, as long as dividend policies stabilize and the dividend–price ratio resumes its mean-reversion, though perhaps to a new mean. (An Appendix, available on request, provides an example involving repurchases.) But there can be problems during transition periods. For example, if the fraction of firms that do not pay dividends steadily increases, the market dividend–price ratio is probably nonstationary; it is likely to decline over time, and the dividend growth model is likely to underestimate the expected stock return.

Fortunately, the earnings growth model is not subject to the problems posed by drift in dividend policy. The earnings growth model provides an estimate of the expected stock return when the earnings–price ratio is stationary. And as discussed earlier, the model provides an estimate of the average expected return during the sample period when there are permanent shifts in the expected value of Y_t/P_t, as long as the ratio mean-reverts within regimes.

The earnings growth model is not, however, clearly superior to the dividend growth model. The standard deviation of annual earnings growth rates for 1951 to 2000 (13.79 percent, versus 5.09 percent for dividends) is similar to that of capital gains (16.77 percent), so much of the precision advantage of using fundamentals to estimate the expected stock return is lost. We see next that the dividend growth model has an advantage over the earnings growth model and the average stock return if the goal is to estimate the long-term expected growth of wealth.

B. The Investment Horizon

The return concept in discrete time asset pricing models is a one-period simple return, and our empirical work focuses on the one-year return. But many, if not most, investors are concerned with long-term returns, that is, terminal wealth over a long holding period. Do the advantages and disadvantages of different expected return estimates depend on the return horizon? This section addresses this question.

B.1. The Expected Annual Simple Return

There is downward bias in the estimates of the expected annual simple return from the dividend and earnings growth models—the result of a variance effect. The expected value of the dividend growth estimate of the expected return, for example, is the expected value of the dividend yield plus the expected value of the annual simple dividend growth rate. The expected annual simple return is the expected value of the dividend yield plus the expected annual simple rate of capital gain. If the dividend–price ratio is stationary, the compound rate of capital gain converges to the compound dividend growth rate as the sample period increases. But because the dividend growth rate is less volatile than the rate of capital gain, the expected simple dividend growth rate is less than the expected simple rate of capital gain.

The standard deviation of the annual simple rate of capital gain for 1951 to 2000 is 3.29 times the standard deviation of the annual dividend growth rate (Table I). The resulting downward bias of the average dividend growth rate as an estimate of the expected annual simple rate of capital gain is roughly 1.28 percent per year (half the difference between the variances of the two growth rates). Corrected for this bias, the dividend growth estimate of the equity premium in the simple returns of 1951 to 2000 rises from 2.55 to 3.83 percent (Table IV), which is still far below the estimate from the average return, 7.43 percent. Since the earnings growth rate and the annual rate of capital gain have similar standard deviations for 1951 to 2000,

Table IV

Estimates of the Real Equity Premium in Simple Annual and Long-term Returns: 1951 to 2000

The inflation rate for year t is $Inf_t = L_t/L_{t-1}$, where L_t is the price level at the end of year t. The real return for year t on six-month (three-month for the year 2000) commercial paper (rolled over at midyear) is F_t. The nominal value of the S&P index at the end of year t is p_t. Nominal S&P dividends and earnings for year t are d_t and y_t. Real rates of growth of dividends, earnings, and the stock price are $GD_t = (d_t/d_{t-1})*(L_{t-1}/L_t) - 1$, $GY_t = (y_t/y_{t-1})*(L_{t-1}/L_t) - 1$, and $GP_t = (p_t/p_{t-1})*(L_{t-1}/L_t) - 1$. The real dividend yield is $D_t/P_{t-1} = (d_t/p_{t-1})*(L_{t-1}/L_t)$. The dividend growth estimate of the real S&P return for t is $RD_t = D_t/P_{t-1} + GD_t$, the earnings growth estimate is $RY_t = D_t/P_{t-1} + GY_t$, and R_t is the realized real S&P return. The dividend and earnings growth estimates of the real equity premium for year t are $RXD_t = RD_t - F_t$ and $RXY_t = RY_t - F_t$, and $RX_t = R_t - F_t$ is the real equity premium from the realized real return. The average values of the equity premium estimates are $A(RXD_t)$, $A(RXY_t)$, and $A(RX_t)$. The first column of the table shows unadjusted estimates of the annual simple equity premium. The second column shows bias-adjusted estimates of the annual premium. The bias adjustment is one-half the difference between the variance of the annual rate of capital gain and the variance of either the dividend growth rate or the earnings growth rate. The third column shows bias-adjusted estimates of the expected equity premium relevant if one is interested in the long-term growth rate of wealth. The bias adjustment is one-half the difference between the variance of the annual dividend growth rate and the variance of either the growth rate of earnings or the rate of capital gain. The equity premiums are expressed as percents.

		Bias-adjusted	
	Unadjusted	Annual	Long-term
$A(RXD_t)$	2.55	3.83	2.55
$A(RXY_t)$	4.32	4.78	3.50
$A(RX_t)$	7.43	7.43	6.16

13.79 percent and 16.77 percent (Table I), the bias of the earnings growth estimate of the expected return is smaller (0.46 percent). Corrected for bias, the estimate of the equity premium for 1951 to 2000 from the earnings growth model rises from 4.32 to 4.78 percent (Table IV), which again is far below the 7.43 percent estimate from the average return.

B.2. Long-term Expected Wealth

The (unadjusted) estimate of the expected annual simple return from the dividend growth model is probably the best choice if we are concerned with the long-term expected wealth generated by the market portfolio. The annual dividend growth rates of 1951 to 2000 are essentially unpredictable. If the dividend growth rate is serially uncorrelated, the expected value of the compounded dividend growth rate is the compounded expected simple growth rate:

$$E\left[\prod_{t=1}^{T}(1 + GD_t)\right] = [1 + E(GD)]^T. \tag{5}$$