

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

^aNot rounded, 6.98 percent.

^bNot 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.

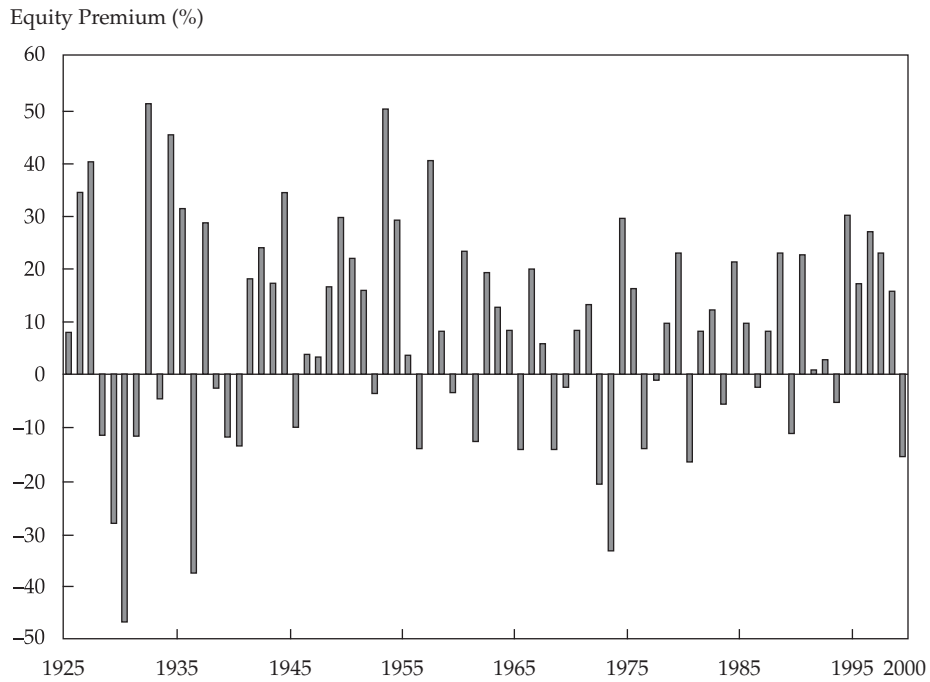
Which of these interpretations of the equity risk premium is relevant for an investment advisor? Clearly, the answer depends on the planning horizon.

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.

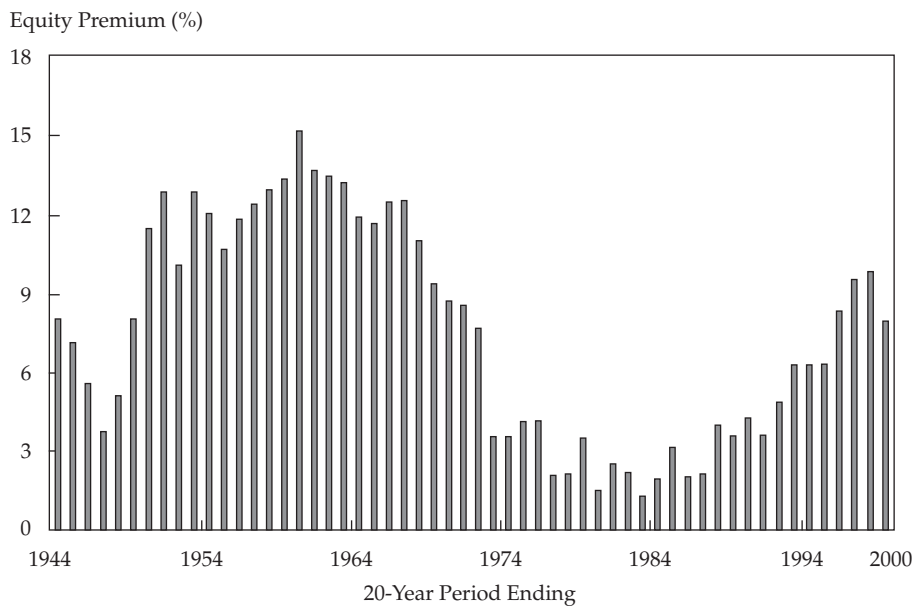
Figure 3 carries out exactly the exercise that Brad Cornell recommended [see the “Historical Results” session]: It looks at stock market value (MV)—that is, the value of all the equity in the United States—as a share of National Income (NI). These series are co-integrated, so when you divide one by the 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



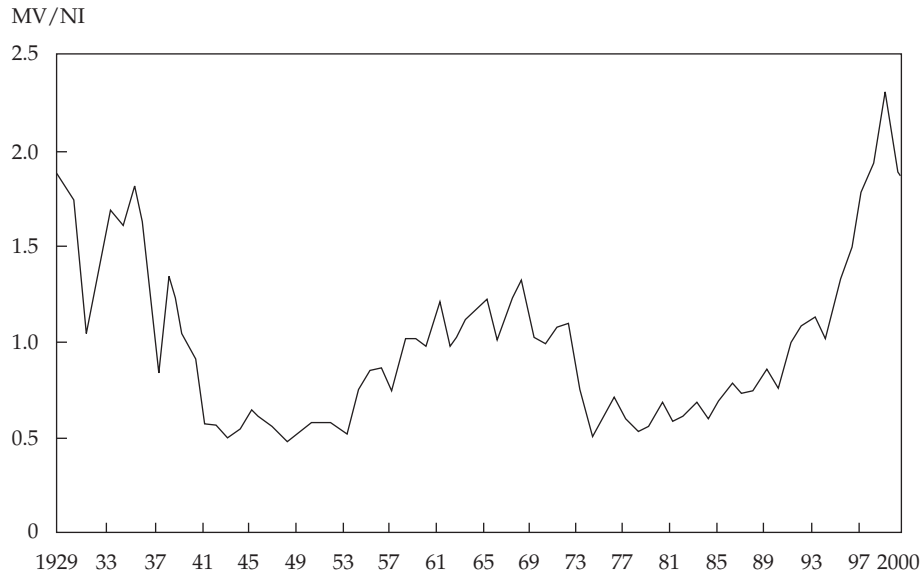
Source: Ibbotson Associates (2001).

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



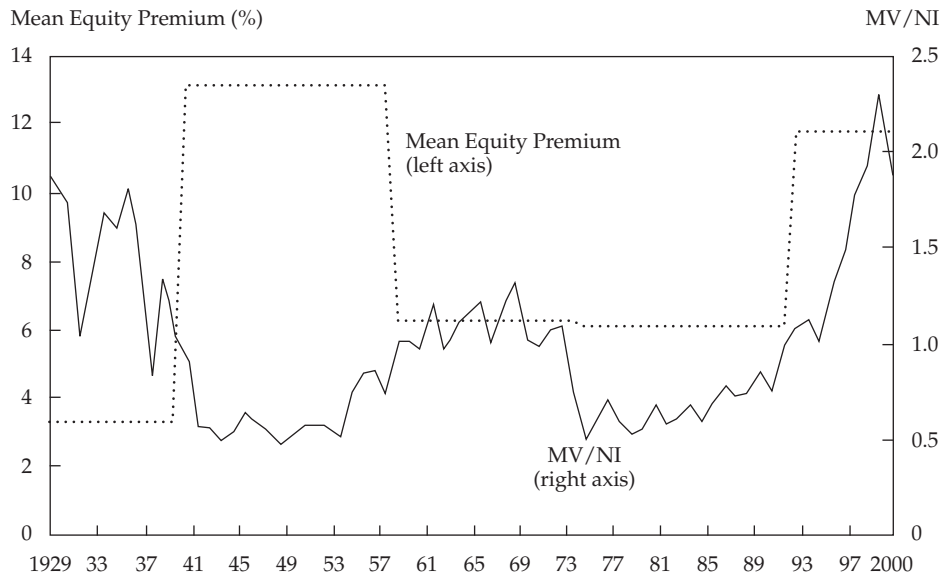
Source: Data updated from Mehra (1998).

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

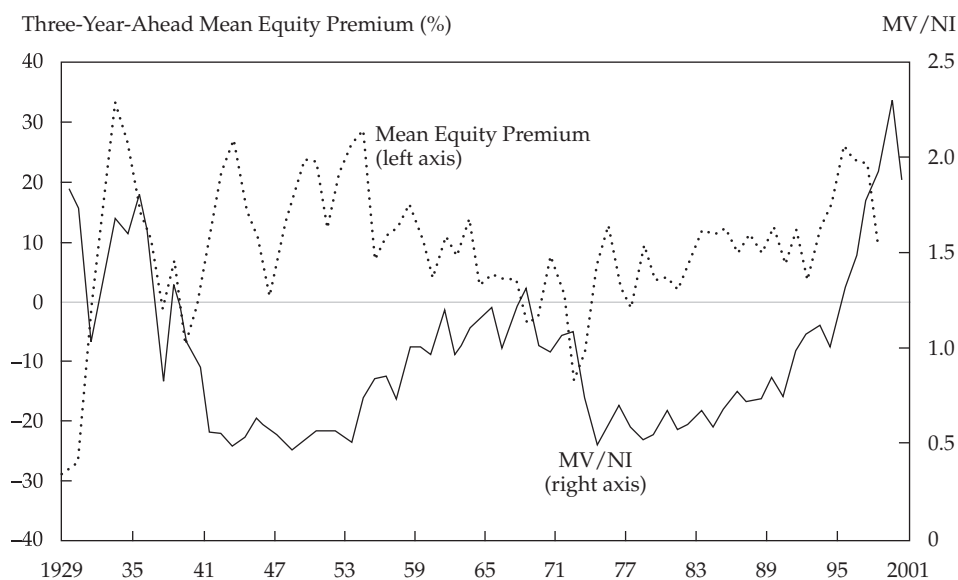
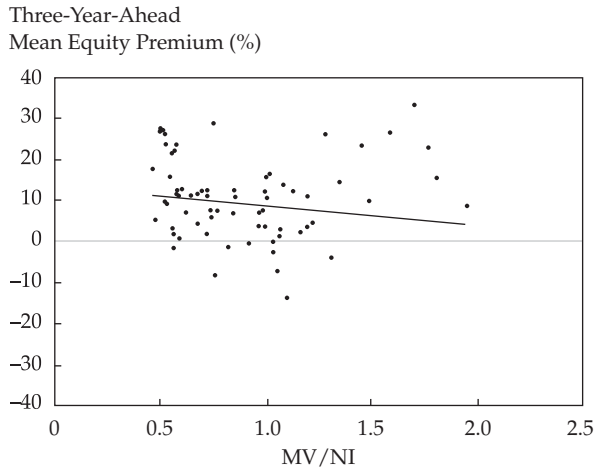


Figure 6. Scatter Diagram: Mean Equity Risk Premium Three Years Ahead versus Market Value/National Income, January 1929–January 2000 Data



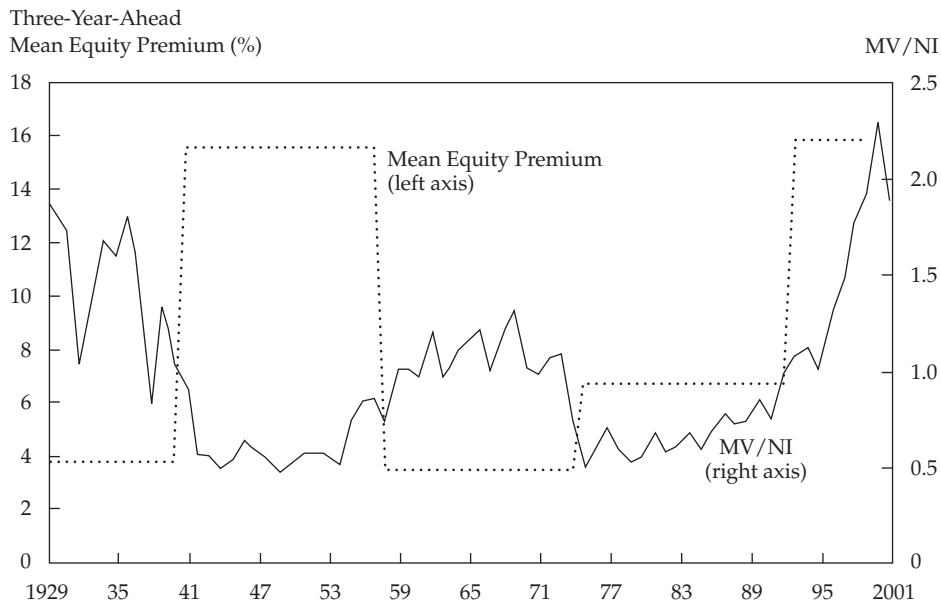
Note: $y = 4.7159x + 13.321$.

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 $MV/Ni < 1$.