

Historical Results: Discussion

Ravi Bansal (Moderator)
Robert Arnott
Clifford Asness
John Campbell
Peng Chen, CFA
Bradford Cornell
William Goetzmann
Campbell Harvey
Roger Ibbotson
Martin Leibowitz
Rajnish Mehra
Thomas Philips
William Reichenstein, CFA
Stephen Ross
Robert Shiller
Jeremy Siegel
Kevin Terhaar, CFA
Richard Thaler

RAVI BANSAL (Moderator)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

LEIBOWITZ: It's exactly linear.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

SIEGEL: Yes.

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

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

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

THALER: Right, right.

ASNESS: You never succeed.

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

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

THALER: Right.

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

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

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

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

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

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

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

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

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

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

IBBOTSON: We'll investigate that!

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

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

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

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

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

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

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