SURVIVORSHIP IN FINANCE

Moderator: R. STEPHEN RADCLIFFE
Speaker: STEPHEN A. ROSS*

MR. R. STEPHEN RADCLIFFE: Stephen Ross holds the Sterling Professorship of Economics and Finance at Yale University. He is also a managing director of WP Capital and is a principal of Roll and Ross Asset Management Corporation. He is the author of more than 50 articles in finance including papers on arbitrage pricing theory, agency theory, and derivatives. He is coauthor of an introductory textbook in finance and has received numerous prizes and awards for his work.

Professor Ross is a Fellow of the Econometric Society and a member of the American Academy of Arts and Sciences. He is also director of the General Re Corporation and of the College Retirement Equities Fund. He is past president of the American Finance Association.

MR. STEPHEN A. ROSS: It's a pleasure and particularly a privilege to have the opportunity to address so many people who are both interested and knowledgeable about topics that are of interest to me. I've had the chance to look over your program; I think it's fascinating. I commend you for it. But I had a somewhat mischievous idea as I was doing so. If I asked, particularly if I were to poll the younger among you, what the onset of modern finance would be, I'd get a somewhat different answer than my own. I suspect that the majority would say, judging from the program, that the onset of modern finance began with the Black-Scholes model and option-pricing theory.

Last week I had the occasion to be with a group of people at an industry conference. They were primarily managers of portfolios and pension-fund managers. I asked them when they thought modern finance began, and they thought it began somewhat earlier than you. They said they thought it began with the capital-asset-pricing models, around the late 1960s and 1970s. I actually trace the roots of the modern subject back a bit further. I traced it to a wonderful, somewhat neglected article in 1937 by Cowles, who examined what we now call the efficiency of markets. And that's really part of what I'm going to be saying. I'm going to be talking about statistics in finance, and in particular, what I think are some of the problems and mistakes that we tend to make when we look at financial statistics.

Efficient market theory lay dormant after Cowles until around the 1950s, and then it picked up in steam in the 1960s and 1970s. It is the empirical basis for what we think of as modern finance. If you look closely, lurking in the background of option-pricing theory, asset-pricing models, and all of the paraphernalia of modern finance, are the fundamental intuitions of efficient market theory. Efficient market theory, at least as it was understood back then, was really a theory about how prices of marketed financial assets have moved over time. The basic argument was really quite a simple one. Of course, all things cease to be simple when you think about

---

them long enough. But at the time, the thought was that the current price was really some sense of the reflection of the consensus of all of the participants in the market. As such, it incorporates all of the information that people have. The implication of that statement is that the movement in price over time has to be random. In particular, the change in price in April will be uncorrelated with the change in price in March; and uncorrelated, for that matter, with the change in price in February or January or any previous date. In what I think of as its most religious form, the efficient market theory says that whatever happens to the price tomorrow is absolutely and completely unrelated to anything that's ever happened in the past.

Well, it's a strikingly strong proposition, particularly from the perspective of modern economics, which tends to say things such as, I think GNP might go up tomorrow, and then again it might not. And if it does, it might have an effect on interest rates, and then again it might not. The efficient market theory says very specifically that, if you were to look at statistical relations between the change in the price this month of some stock, or the stock market as a whole, or bond prices, you won't find any statistical relation between that change in the price and past changes.

It's appropriate we're here in Orlando. I think that one of the most interesting, and one of the most craftsmanlike of the works ever done on this subject was actually done in the 1980s by a friend of mine, Richard Roll. He studied the orange futures market, and the orange futures market is actually centered in Orlando, Florida. Approximately 90% of the oranges, at least domestically, that can be grown for frozen concentrate come from around here. As a consequence, the federal government and, to some degree, the state government have spent literally billions of dollars on weather forecasting in this area. Satellites are tuned to stations located on the ground here. The largest concentration of weather forecasting capability anywhere to be found in the world is all focused right here on Disney World. This is all designed to predict what the minimum temperature is going to be tonight. If it gets too cold tonight, then the orange crop will freeze, and if the crops freeze, orange prices will shoot through the ceiling. There is actually something you can do about this if you're a grower. You can go out and put smudge pots or fans out, and you can do something to alleviate the problem.

The National Weather Forecasting Service has rolling forecasts of what the minimum temperature is going to be tonight. At 7:00 a.m., it produced a forecast of what the minimum temperature will be tonight, Thursday night, and Friday night. And that is updated during the day. So if you think about Friday night, starting with a 7:00 a.m. forecast this morning, you're going to hear perhaps half a dozen significant forecasts of what that minimum temperature will be, culminating with a 7:00 p.m. forecast on Friday night.

Well, the orange futures market is, of course, following all these predictions, and the price of orange futures contracts are changing over time. Dick Roll studied whether those prices really were related, for example, to changes in the weather forecast. He found that, in fact, the change in the orange futures price for any day was absolutely unrelated to past weather forecasting changes. It obviously changed as the weather forecast changed. But for example, the Wednesday price change was uncorrelated with the price change from Monday to Tuesday, i.e., the futures price was uncorrelated over time.

But perhaps most interestingly, the futures market closes at 3:00 p.m. The last forecast of the minimum temperature is done at 7:00 p.m. Dick Roll discovered that if the weather forecasting bureau used the 3:00 p.m. closing futures price, it could improve the forecast of the minimum temperature. I considered that to be just an absolutely wonderful tribute to the efficient market theory, but it wasn't greeted with much enthusiasm by the weather forecasters.

The weather forecasters are not the only ones who fail to view this theory with enthusiasm, which I consider so powerful and wonderful, at least in the social sciences. This has been a controversial theory from its birth. The first studies of this were really studies of price movement over time, that was predicted by the theory. I describe the studies in random fashion. But other studies, which I actually found more compelling, are going to be the focus of part of my talk. Those were studies of performance of professionally managed mutual funds.

A typical study went something like this. You took all the funds you could find. In January 1970, there were about 90 funds. You looked at these funds and asked, how did those professionally managed funds perform in 1970? And the answer came through with daunting clarity. There was, of course, a distribution of these funds; some were doing very well and some did very poorly. The average fund underperformed the market—the market at the time was taken to be the Standard & Poor's (S&P) index—and the average fund underperformed the S&P index by a small amount that seemed uncomfortably close to precisely average fees changed by funds. So if you looked at the mean of that distribution, it was down by about 50 or 60 basis points from what the index was doing. This was not greeted with enthusiasm by our friends in the investment management community.

A good friend of mine, Paul Cootner, who has since passed away, was one of the pioneers in this. He was involved in what is now a famous exchange in this area. He was giving one of the first talks to an industry group explaining this result. Someone from the industry arose and gave a familiar trite expression and complaint. He said, "Well Professor, that's nice to hear, but if you're so smart, why aren't you rich?" And Paul, who had a wonderfully dry wit, got up and said, "I understand but tell me something, if you're so rich, why aren't you smart?"

I think this might have gone too far. All the studies I'm describing of the performance of these mutual funds have now been updated. And I'm going to tell you something about what the new results in this area are. As I said, the most compelling things for me about those industry studies showed that the average professional manager just about did the market minus the fees charged. Industry professionals, though, suggested an interesting second test, which is actually the focus of what I'm concerned about. They said, "Yes, it's fine to say that the average underperforms, but like all industries, there are good folks and bad folks. What's really critical is not so much whether the average professional manager beats the market, rather what's really critical is to recognize that there are good managers and bad managers. Some can outperform and some can't." And to test this, the natural way to think about it was to see whether performance was persistent over time.

The standard tests from the early 1970s were tests of the following sort. You would look at the performance of managers in 1970 and rank them from top to bottom.

Say there were 100 managers; you ranked them from the best to the worst. And then you'd see how they did in 1971, and again rank them from the best to worst, and see the relationship between the ranking. If you were in the top 10% in 1970, was it likely that you were in the top 10% in 1971, or at least above the average in 1971? Well, the news came back just like the previous news. There was no correlation. And we don't even have to talk about statistically significant correlation. There really wasn't any correlation at all between the rankings of managers in 1970 and 1971, or for that matter in any other years the people did the study. The manager at the bottom of the totem pole in 1970 had the same probability of being at the top in 1971 as the top manager in 1970, and that seemed to be true if you held for several years. Well, needless to say, this endeared the academic community further to the industry.

But as I said, these studies have all been redone. And there's been a spate of new studies in the late 1980s and the early 1990s resuscitating this effort. These studies have come to really quite markedly different conclusions.

First, these studies have discovered, and this is going all the way back to the 1970s up to the present, that the average performance is actually better for professional managers than for the market as a whole. And second, good performance seems to persist. So managers who beat the benchmark in one year beat it in another year. Let me give you an example of this.

We look to the sample of growth-fund managers. The results are not sensitive to what time period you look at, and they're not sensitive to how you actually measure their performance, and they are not sensitive to whether you use growth funds, or value funds, or all the funds for that matter. The performance of growth funds is being measured to the median of the sample of these managers. There are 181 funds in all. This is a real-world study. I wish there were 180, so I could say 90 were above and 90 were below. I can't. So I have to toss the one somewhere. The median is around the 90th or 91st fund. And we looked at the alpha, which is just one measure of the excess performance for these funds. Well, how did they do?

There were 91 winners in 1984–85. There were 90 losers in 1984–85. I tossed the extra one into the winners' spot.

So, for example, of the 91 winners in 1984–85, 58 of them continued to be winners. That is, they beat the median of their sample in 1986–87. And only 33 of them stumbled back to lose to the median in 1986–87. And similarly, of the 1984–85 losers, 57 actually continued to lose in 1986–87. This is very strong and statistically significant evidence of persistence of performance. I suppose the most persistent you could find would say that every winner stayed a winner, and every loser stayed a loser. That's much too much.

Using the cross-product ratio test; 3.04 is just a ratio of the product of that 58 and 57 divided by 33 times 33. That shows that 0.2 of 1% is the chance that this could have occurred just by chance alone, if in fact, there wasn't real persistence in the data. Well, this was the typical kind of study that I was confronted with in the late 1980s and 1990s. And this was troublesome to me. I asked people why their results were so different from the results that people had before? The answers I got

4

were interesting. The answers were of the form that in the 1960–70s, so they really didn't know a lot about conducting these studies. Statistics have improved since then. That didn't seem a very satisfying answer to me.

It's not like saying they didn't know a lot back in the 15th century about statistics. And the truth of the matter is, the statistics used in these studies haven't really changed very much. Some of the tests are new. But they wouldn't have changed anything, had you looked at them back in the 1960s and 1970s. More pertinently, it's not like you're talking about people from the 18th century who are dead. Most of these people are alive. They take umbrage at being accused of not knowing statistics and doing their test improperly. I thought perhaps a better explanation was that something was wrong.

Instead of 100 funds we have 1,000 or so funds, or 800 and some. And that was another explanation; now the data are much more voluminous, they go back much further, and they are much more complete. In fact, that's not the issue. The tip-off is that the average manager in this sample is outperforming.

But the average manager outperforming is to me a signal that something else is going on. And what's going on is survivorship bias. I will describe a classic example of a set of alpha tests. Many of them have been run on a data sample that is inherently skewed because you start at the end and look backward. It contains survivors. Indeed, the way most of these samples have been collected is that they would start, say, at the year 1990, and gather the data on all the funds that were around in 1990 by going backward in time. If you go backward in time, say to 1985, you're looking at all the funds in 1985 that survived all the way through to 1990. That's a fairly significant thing to discover. That clearly biases the results upward. And that's sort of the first-blush effect of survivorship.

If you have a surviving sample, and you've weeded out all the funds that somehow or other didn't meet the market tests, then you've weeded out all the funds that didn't perform to the averages. And so it's not at all surprising that the average performance of this surviving sample would exceed the average performance of the market, or any other measure of the broad market return. What is surprising is that you get this kind of persistence. Sometimes this persistence, as I label it, is called hot hands.

It was clear people were prepared to accept that survivorship bias buoys the averages up. But what should survivorship and survivorship bias have to do with them showing persistence? Survivorship bias doesn't tell you they persistently do well among the survivors. And I didn't really have a simple answer to that. But I had long ago learned to respect survivorship bias. Survivorship bias is wonderful. It is ubiquitous. I'll show it to you in many different examples in finance later. And it really lies at the heart of how we understand financial, and for that matter, I think, much economic data.

Suppose you have two funds; lets call them X and Y with manager X and manager Y. I'm going to look at their performances over a ten-year period. Both X and Y survived to the end of the ten years. Manager X beat manager Y in the first five years. Let's say it's 1980–90. So I'm going to give you two pieces of information

about them. I'm going to say that X was better than Y in 1980–85. And both X and Y made it all the way to the end to 1990. What would you deduce from something like that? Well, my first deduction would be if X beat Y in 1980–85, and both X and Y stayed in the sample and they both survived for the ten-year period, Y's performance must have improved, and X didn't have to do as well. Manager X, in effect, could coast in a conditional probability sense. And Y had to improve in that sense. So I expected that survivorship would lead to a flip-flop. And that would make the results on mutual fund persistence even more impressive.

Now I'll tell you something a little bit different about X and Y. Fund Y managed to beat the market by half a percent the first year. It lost by 40 basis points the second, and it lost by 30 basis points the third. In other words, fund Y just hugged the S&P index. Fund X, on the other hand, was taking wild chances. It had a huge standard deviation. One year it was 20% above the market, another year it was 20% below the market. But they both survived all the way to the end. And conditional in surviving, if they both survive all the way to the end, the riskiest fund does the best.

Suppose the two managers X and Y both survive, and one of them takes much bigger bets than the other. That's the one with the broader standard deviation. But to survive they both have to beat a hurdle. Well, conditional on surviving, the expected value for manager X is higher than that for manager Y. If two people go to a casino and one takes huge bets, and the other one takes little bets, and they walk out at the end of the day and both survive, chances are the one who took the bigger bets won more money. Each had the same chance of winning or losing. But given that they both won, the one who took the bigger bets won more. That, in fact, explains the data from the 1980s to 1990s. I can't say it completely explains the phenomena. This is a debate that's now going on. But this survivorship bias in the sample is powerful and it does move in the right direction.

That essential finding, that higher risk funds have to do better, explains why you have persistence in the data. Persistent ones are the higher, riskier ones that survived. That explains the phenomena we're looking at with the new material. That seems to be an interesting result. But it's kind of a negative result. We then went ahead and ran simulations. We did 20,000 simulations, and we cut off the bottom 5% or the bottom 3% of these funds every year to just see if we could duplicate the results in the studies. And, in fact, we could.

Having said all that, it's kind of a nihilistic argument. It doesn't really tell you what to do with it, nor does it really give you a flavor yet of why this is important. But one of the most interesting features of this work for us was that we actually were able to find a nice trick that removed the effective survivorship bias. We used something called the appraisal ratio. Those of you who studied old finance might know about this. It's just a ratio of whatever performance index you're using to the standard deviation, or the breadth of the density. If you rank funds on that basis, you can get rid of the survivorship bias entirely. And, in fact, you replicate the results that people found in the 1970s and 1960s in the earliest studies, which suggests to me that this impact is one of survivorship bias.

Now some people, when confronted with this view, say the study wasn't done right. Why didn't the researchers just go back and do the study right? Why don't they just

look at all the funds in 1980 and follow them forward, instead of looking at all the surviving funds in 1990? That's a nice answer. But it's not a satisfactory answer. And it's fundamentally not satisfying for finance and financial economics, and economics, in general, and I suspect even for the social sciences. The very nature of what we do in finance means that we generate huge amounts of data, which we throw away. We only become interested in things after we've thrown out the data.

For those of you trained as physicists, if it wasn't that way, you would have to carry the lifeline of all the objects in the world. The world would just grow at some incredibly fast exponential rate. We'd be drowning in our paper. That's not the way the world is. When a company dies, all of its accounting material is thrown away. We don't keep the old records of old companies. We don't keep the old records of old funds or of old stocks. We just throw that away. We have no choice. Unfortunately, 25 years later, someone says it would be interesting to see how the performance of the stock was correlated with earnings' announcements.

We live here today and with some exceptions but not generally, we just have to live with what we have today, knowing that what we have today survives. So I think the game is good statistics; not by going back trying to gather up information that isn't there, but rather recognizing the survivorship and understanding its implications. I think this is critical, and I think it's actually inherent in what it is we do. It separates some of the statistics we do from the statistics in other sciences.

Let me give you some examples of survivorship bias that we've been working on and how it impacts the way we think about statistics. One of the reasons I'm so attracted to all of this, and in particular to this kind of bias in data is because I find it ubiquitous. I see it everywhere.

Second, the mathematics that lie behind the computation and the analysis of these phenomena are precisely the same as the mathematics you're going to be doing in the next couple of days. Option-pricing-theory mathematics are exactly the same as the mathematics here. After all, if something survives, it means it got above some sort of a hurdle. The mathematics of its performance, conditional on getting above the hurdle, are exactly the same as the mathematics of the stock price, conditional on getting above some exercise value. In other words, the mathematics of figuring out the value of a call option are exactly the same as the mathematics of figuring out the statistics that I'm concerned with here.

Interestingly, both literatures have developed in parallel, and now we see them as having a symmetry. We found some interesting results in one area that influence the other. And that seems to be a two-way road. Problems that are solved in one actually have implications for the other. And intuitions from the one have intuitions for the other. So one of the things I'm working on is the intuition of survivorship bias brought back to option-pricing theory.

Let me give you some other examples. How many of you have heard of, invested in, or associated with at any time an emerging market fund? Have you heard of the emerging market? The hoopla about them is based on the observation that Hong Kong and Singapore and Malaysia would have been wonderful countries in which to invest. Or Germany would have been a wonderful country in which to invest, back

when it didn't have the market. And of course, the premium one is Japan after World War II; what a wonderful market! As a consequence, people are setting up emerging market funds to invest in all sorts of little markets around the world. I think this is interesting. But I don't call it emerging markets; I call it *emerged* markets. I would like someone to tell me that such and such a market is emerging. That would be particularly pertinent information. All they can tell me is that a market has emerged.

No one tells me that they might have invested in Peru in 1971 or in Bolivia in the 1960s. No one tells me about the emerging markets that didn't emerge. They only tell me about the emerging markets that did. That is the essence of survivorship bias. And I suspect that survivorship bias, when looked closely at, and I haven't looked at it, would have a profound impact on how we think about these data. And after all, to point at Hong Kong and say it had wonderful double-digit returns, 25–30% returns for a terribly long period of time, tells me nothing. All it does is look at a single market among all the markets you could have looked at, and picks the best one. There are many examples of this. And I think it's time to think seriously about not being too beguiled by that.

Here's another example that has nothing to do with finance. We've been looking at the data on the Nile River. Wonderful, old statistical studies have been done on the Nile River, because we have over 4,000 years of data. For 4,000 years we've collected the high-water mark on the Nile River. It's a grand data series, and several statisticians have written about it. The theme that emerges from this is interesting. There are apparently cycles that one can observe in natural phenomena. So if one looks at the high-water mark, one finds that it tends to have broad, long cycles. It rises for long periods and then falls for long periods. That was considered to be something associated with fundamentals of geology, or the mythology of rivers, or what have you. I don't think so.

What is the one thing we know about the Nile River? We know that it has survived for 4,000 years. Rivers don't survive for long periods of time. How do rivers die? Well, there are two ways that rivers die. They either dry up or they flood over and they become lakes. The latter has happened in the Midwest this past year. Tributaries of the Mississippi have became lakes. So we know that the Nile never did that. It has never dried up in 4,000 years, and it never became a lake. There's only one thing it must have done. It must have cycled between being so low that it would have dried up or being so high that it would have become a lake. There was no choice. I don't think this has anything to do with the natural phenomena of rivers, that the high-water mark fluctuates over time. Leaving aside the biblical explanation for these cycles, it seems to me you're just looking at a river that lasted for 4,000 years. You are looking at something that has lasted for a long time (which is actually quite short by geological time as rivers go, although it's long for a river and short for a continent).

Nonetheless, it's our suspicion that that explains what happened in these data. There's a financial analog. People look at long-run market studies and find what's called reversion to the mean of cycles. In the 60 years of U.S. data, it's tough to find a meaningful cycle. By a cycle, I mean that if the market outperforms some sort of a price/earnings benchmark, or something like that for a long period of time, then it

tends to come down. Or if it underperforms for many years, it tends to come up. It's difficult to bet money on these cycles. I think back most recently to the end of 1989. There were newspaper and financial articles about how the 1980s will never be repeated again and people should pull out of stocks. Well, the early 1990s were not all that bad for the stock market.

Nonetheless, there's a nice intuition about that; investors are faddish, they go a little too far in one direction or another. It's tough to find this result in the U.S. data, but you do find it when you look back 250 years at the British data. You find it in the U.S. data if you go back 150 years. Some very old companies were traded in the late 18th century in this country. Aetna Insurance is one of the oldest companies ever traded. And if you go back and look, you'll find in those data these kind of Nile-like cycles. It goes up, then sort of predictably comes down, and predictably goes up again. But I think that's just because we're looking at surviving markets.

You don't find that, by the way, if you look at the German market. The German market's been here and gone. It disappeared in two wars; it disappeared once because of hyperinflation, it just got so high so fast that the bubble burst because it was just printing money willy-nilly, and then that market disappeared. And then it disappeared again when the country was destroyed in the war. A stock market dies in at least two ways. It either goes up too fast, in which case the country is having incredible monetary problems and the market is simply reflecting that, and then there's a bust and the market stops, or it goes down because the capital stock of the country gets destroyed in a war or for some other reasons.

So when looking at England, that lasted for 300 years in its market; you're just looking at the Nile River, written in financial terms. That's a market that lasted for a long time. The only thing it could have done was to have cycles around its trend. It never went up too fast, and it never went down too fast. If it doesn't go up too fast or down too fast, it has to cycle; all of which takes me back to the moral of my tale.

Efficient markets are really to finance what Newtonian mechanics is to physics. We now play around at the fringes of it, developing some quantum theory. We're working on that now in the information-based theories. But to a very powerful first order of approximation, unless you mislead yourself by looking at the statistics, you're going to find that the markets are efficient.

On the other hand, it's extremely beguiling to look at the statistics. If we look at the sky long enough, we see bears and cats in the clouds. You look long enough and you see anything you want to see in these data. A friend of mine, a Turkish economist, said that if you torture the data long enough, it will confess to any crime.

Our shop is now being brought all sorts of portfolios that I consider prime examples of misunderstood statistics in finance. These are mortgage portfolios that have just blown up and have been destroyed. I have a son in college, and my son is a history major. But he was reading the paper and he saw the term *financial engineering*. He asked me if I was a financial engineer. And I thought about it some, and I said, "No. I consider myself to be a financial pathologist. I practice forensic finance."

FROM THE FLOOR: Did you study any persistence in the nature of the distribution of returns? Based on what you were saying, I would expect a certain pattern of distributions among the ones that performed, and a different one. I'm wondering if you did any studies of consistency of distributions.

MR. ROSS: I think that's an excellent question. The answer is, we did not. We focused on the first and second moments for this study. So we did look at the mean obviously. We also looked at the persistence of the standard deviation. And what we found was precisely what I had described. That cross-sectionally, the funds that were the highest standard deviation were, in fact, the funds that had the highest return because they had survived. So we didn't look beyond that at other moments.

MR. STEVEN P. MILLER: Are you currently gathering data to do a study and including those that may not survive to the end of your study? And do you think that the CAPM relationship between risk and reward is just a belief in statistics with a survival bias?

MR. ROSS: Both of those are good questions. For the first question, the answer is no. I'm not gathering any data. I have graduate students who are; they're presumably working for my coauthors. I've kind of lost interest in that specific question because there doesn't seem to be anything conceptual about it that I'd like to learn more about. And I'm tired of waiting to see what the data have to say.

The second question is about the capital-asset-pricing model. I've given some thought to that. And what is striking is how little explanatory power the model seems to have, even with survivorship bias. So that is something I'm thinking about. But I don't have any definitive, concrete answers for you now.

MR. JOSEPH J. BUFF: I'm curious whether you thought at all about applying some of this analysis to other specialties of managing investments. For instance, regarding the insurance industry and the relative performance or skills of insurance agents, presumably to some degree they correlate from period to period. But there's also a survivorship issue clearly in the insurance industry. Perhaps it's an idealistic perspective, because again it's not clear what that would really tell anybody if we knew it. Have any studies like this been applied to things other than to the management of investments?

MR. ROSS: I am not aware of any similar studies. I can tell you something from the investment world, though, that is similar, but that is not really about investments. In the large investment banking houses, traders tend to elevate to the top positions, at least in terms of compensation. There's something called a trader's option. You take a lot of risk, and maybe you are fired if you don't work out, and you move to another company, or, on the other hand, if you do work out, you make a great deal of money. And that's called the option that the trader has. The problem is that whenever people watch a trader, they say how well the trader did this year; or I suspect the insurance agent, or someone like that. They say, "Look how well that person did this year." They lose track of the fact that the question isn't just how well the person did against the average, but how well he or she did conditionally, on being better than the average. So I've been thinking about applying this kind of analysis to examining traders and what is really significant about a trader's performance. Quite

typically, people will show you a list of 50 traders in a very large organization. You look at the person at the top and say, "Look at all the money that person made. That person is really good." Well, the one thing we know is that the very fast person in the race had two things going for him or her. He or she was fast and lucky. And you win races by being both fast and lucky. You don't win them just by being fast. I want to remove that component of luck, and get a more sensible, statistical hurdle for what is meaningful in terms of additional performance.

FROM THE FLOOR: You mentioned that stock markets are driven by monetary policy. And I wondered if you could apply the survivorship metaphor to looking at interest-rate movements and the yield curve.

MR. ROSS: Yes, that's part of that interesting relation between option-pricing theory and what we think of as continuous time finance on the one hand and this survivorship thing. The early studies of the models of the term structure of interest rates all focused on signing some nice, attractive probability distribution for the movement of the interest rate. That would enable you to solve for the value of financial claims bonds and things like that. It turns out that survivorship presents you with some very interesting ones. Look at all the interest rates that are generated just by an absolute random walk. And then say that interest rates can't go negative. So just slice away all the parts that could ever go negative. And look at the ones that survive. If you look at the interest rates that never go negative, if you use that for your model for interest rates, you can solve some problems, and value bonds and things that aren't solvable by other techniques now. That's a subject of interest to me.

MR. CHARLES E. MOES, JR.: I enjoyed your comments about whether the professionally managed funds outperform the broader market averages. But for me, personally, the question is, load or no load. I guess a lot depends on when people are trying to sell me something, in terms of where they're coming from. Have you done any analysis of which is more effective in this regard?

MR. ROSS: I have done absolutely no analysis, and I never buy a load fund.

MR. ROBERT J. JOHANSEN: I'm chairperson of a project oversight group investigating the investment returns of companies with significant investments in junk bonds. We were able to trace the period from 1986 to 1992. The preliminary results will be presented in a session at this meeting to which you are all invited. One of the problems we have is that of survivorship. There were some very well-known, let's say, decedents, among our initial group of junk-bond companies. We have essentially traced the rates of return with and without unrealized capital gains. And we've also traced such things as the rate of growth of these companies and some other statistics. So it's a statistical study. But the problem is survivorship. And we have traced these returns year by year. But, of course, the groups of companies have changed. And frankly, we haven't found a good way to reflect this in our study except to say that some of these companies did not survive. And we compared them with our matched set of companies that did not have significant junk-bond investments. And one of our control companies did not survive. But the numbers of nonsurvivors or decedents is very strong. The problem is, we have these results traced through 1992. But we'd love to have something. Suppose these companies

had survived; one in particular. And I will invite people to come to our session and tell us what they think we ought to do.

MR. ROSS: Because the vast majority of these problems that show up empirically have no analytically tractable results, the most successful technique we have found is bootstrapping. We take the null hypothesis, whatever that might be, namely that the returns are the same for these funds and others, or whatever, and then we just generate thousands of things. Computers are wonderful. We generated 20,000 of these for our simulations here. I gave this talk at Stanford, and I was asked why I did 20,000. Someone asked, "Can you prove to us that 20,000 is enough?" I said, "I did 20,000 because I didn't want to have a discussion about whether it was enough. Twenty thousand is enough." It's a nice, simple technique. You just generate tens of thousands of these replications for your sample and see where the actual one you have falls in the distribution under the null hypothesis. And that turns out to be the most powerful technique we have for dealing with these issues.

MR. JOHANSEN: One of our problems is a small number. We took all of the sizable junk-bond companies, and all of the sizable junk-bond companies numbered 63. That's what we started with. And so the results will be an analysis of a small number of companies. But that's our total universe.

FROM THE FLOOR: In view of your very interesting comments about survivorship of markets, in particular the trials and tribulations of Germany, do you have any comments on trends and the efficiency of currency markets in survivorship bias in that market?

MR. ROSS: I actually have a graduate student who's working on currency market. There seems to be the same source of elements of survivorship bias there. One way out of survivorship bias, at least in some of the phenomena we're looking at, is to look at smaller and smaller time periods. So we actually have in those markets the ability to look not merely daily, but throughout the day. And so you can get a lot of information. It comes to be the same thing. If you look at a very tiny fragment of a surviving piece of data, it really does look a lot like one that doesn't have the survivorship bias in it. Survivorship bias seems to be with the large but not with the small, not locally. So currency markets have the opportunity to avoid that. On the broadest scale, they have survivorship along with the countries that they represent.

FROM THE FLOOR: The problems with these sorts of performance studies are surely also bound out with East Europe, the differences of investment style within the sample. So there are clearly going to be some correlations. I would have thought some of the persistence is attributable to that fact. I wonder if you have a comment on that.

MR. ROSS: In fact, that is the case. There is correlation. The sample is not cross-sectionally independent. By defining by style you eliminate a lot of the correlation.

FROM THE FLOOR: But some style still remains no doubt.

MR. ROSS: It does. But we actually were able to sort of factor in little cuts of the data and eliminate most of the cross-sectional covariance.

FROM THE FLOOR: Right. The conclusions are certainly rather nihilistic. Do you have any comment about a strategy for selecting managers that's consistent with your interpretation of the data? I would argue it is probably a no-hands strategy, that's to say, passive management.

MR. ROSS: No, I think the message I come away with is that I can fix the statistics. The performance analysis of any manager is not going to be statistically significant over the lifetime of any manager. That includes Warren Buffet. So it seems to me that you have to say, "I have a harder problem at hand than I thought I did." Creating good performance management is not unlike creating widgets in a factory. And you don't necessarily judge it by just looking at the widget. You go into the factory and say, "Well, what are the inputs in this production process? How do they go about doing it?"

Say I had two managers, one of whom outperformed the other in a five-year period, and both of them outperformed the market. The one who outperformed by the least amount happened to work 50-hour weeks, studying and visiting companies, and doing all the things that you think the manager ought to be doing. And the one who won lived in Nassau and did all his work in December when time came to give the annual report. Knowing something about the process would tell you something about which one you might prefer to have.

MR. RADCLIFFE: Let me take a moderator's prerogative here, and just ask one final question, which is a little off the point. I noticed you had derivatives in your studies. And the recent volatility in the bond market has been blamed on derivatives. I wondered if you could comment on this villainy. And how much potential is left in the current market, or is this mischief to persist?

MR. ROSS: I actually think we've misidentified the villain. I'm in the camp that says derivatives are good rather than bad. Like everything, there are good derivatives and bad derivatives. But to tar and feather the entire industry for having caused volatility, and being in consternation about what we see, to me is just bizarre. This began in 1987, with all the breast-beating about the crash in the stock market, and how this was caused by portfolio insurance. Detailed, cross-country studies were unable to find any relation. Now if you think that this person is a villain, and you see him or her holding the gun, and you see the victim dead on the ground, but the person is holding a 45 and the victim died from a 38, it's kind of hard to pin it on your suspect. But nonetheless we persist in saying that derivatives are the cause of the volatility. The markets have changed dramatically, and so we think that those changes in market structure must be associated with any increments in volatility. From a statistical perspective these increments in volatility we see are absolutely explicable, and they look exactly like the increments we saw prior to having any kind of a derivatives market. So I encourage you rather than discourage you to look at derivatives.