RECORD, Volume 22, No. 3*

Orlando Annual Meeting
October 27-30, 1996

Session 12TS
Forecasting Social Security Actuarial Assumptions

Track: Education and Research
Key words: Education, Research

Moderator: KRZYSZTOF M. OSTASZEWSKI
Instructors: RICHARD S. FOSTER
EDWARD (JED) W. FREES
HOWARD YOUNG
Recorder: KRZYSZTOF M. OSTASZEWSKI

Summary: This session presents a forecasting model of economic assumptions that are inputs to projections of the Social Security system. Social Security projections are made to help policymakers understand the financial stability of the system. Because system income and expenditures are subject to changes of law, they are controllable and are not readily amenable to forecasting techniques. Hence, we focus directly on the four major economic assumptions to the systems: the inflation rate, investment returns, wage rate, and unemployment rate. The other major input to Social Security projections, population models, requires special demographic techniques and is not addressed here.

Mr. Krzysztof M. Ostaszewski: Our first speaker is Jed Frees. He is a Fellow of the Society of Actuaries (FSA) and the Time Insurance Professor of Actuarial Science at the University of Wisconsin. He will present a paper that he has written with several co-authors titled, "Forecasting Social Security Actuarial Assumptions." Our next speaker will be Howard Young, who is also an FSA, and currently a professor at the University of Michigan. He chaired the Technical Panel on Assumptions and Methods for the current Security Advisory Council. Our last speaker will be Richard Foster, who is also an FSA and holds the position of the chief actuary of the Health Care Financing Administration. Finally, I'm an Associate of the Society of Actuaries (ASA), a professor and the director of the actuarial Program at the University of Louisville.

Let me begin with a couple of comments on the subject of the session. The topic of Social Security and its long-term financing problems is now becoming a hotly debated one—maybe not as hotly debated as Medicare financing, but it is coming into the mainstream. More and more often we hear about the so-called "Old Age Crisis." From my perspective, I would like to suggest that Social Security's function within the financial system of the nation, in a sense, is the provision of capital assets to that part of the society that generally tends not to have them. If you study statistics of wealth distribution in the U.S., it is astounding how Social Security wealth, as a financial asset, dominates the wealth of a majority of the people. That is probably the major reason why it is a very popular program, and why it is a program that plays such an important social role in the nation.

We often speak about the problem that we have with the current long-term financing, and the actuarial deficit, as the "old-age crisis," because of demographics. I would like to make a suggestion that this is not really an old-age crisis, but a crisis of pricing of capital assets. A large portion of the population, marginal decision-makers (decision-makers whose decisions have influenced at least some prices of capital assets), are not really seeing the economic consequences of their decisions. I used to think that the baby boomers will have to, at some point, start saving because they will not be able to replace the standard of living with the savings they have right now, but the savings rate seems to have dipped again in the U.S. Now, I'm more and more convinced that what is happening is that the baby boomers are really playing chicken with the government, and they're waiting to see who blinks first. I don't know who will, but I must say that I'm very concerned because, as I have often said in discussions of this issue, I think it is very important that people understand that they need to feel their pain in order to make the right decisions.

At this point, I will give the floor to Jed Frees to speak about forecasting Social Security actuarial assumptions.

Mr. Edward (Jed) W. Frees: This is advertised to be a teaching session with no experience necessary. My role is going to be to introduce people to the notion of how to use statistical techniques in an important problem. I'm going to show you a lot of pictures to demonstrate how to forecast things. I'm fortunate to have a number of expert panelists here who are going to talk about this important problem. We heard a couple of different times already about how the Society of Actuaries feels that Social Security is important. One of the things that I want to try to convince you of is that the techniques are not just important for Social Security, but for other kinds of forecasting—for all macroeconomic problems.

The background of this work is found in a paper that I've written with a team at the University of Wisconsin, my two colleagues, Margie Rosenberg and Jenny Young, as well as a couple of former doctoral students, Yueh-Chuan Kung and Gary Lai. This work was sponsored by the Social Security Administration.

There will be five portions to my talk. I'll speak about the motivation, assuming people here are not experts in Social Security, but at least everybody is affected by it. Then I'll talk a lot about the data and then I'll go over the models and the forecasting. I'll be using extensions of techniques that you might have seen on the Society's actuarial examinations syllabus, Course 120 and Course 121. My talk will conclude with a sample investment fund.

The particular portion of Social Security that is of interest is Old Age Survivors and Disability Insurance (OASDI). It's a piece that affects probably everybody in the room. Almost every western country does have a Social Security system, so the kind of problems that we're experiencing here are not dissimilar at all to other systems. It's a very old program that started in 1935. It is large in almost every sense. It covers virtually every worker in the U.S. There are big amounts of money involved, and it provides benefits to many people. There's a large amount of money associated with it as well, and at least at the end of 1995, there was about $440 billion invested.

I want to talk about the future here. There are really no new ideas here. Right now we're taking in more money than is going out, but it's projected that in about 24 or 25 years this will not be the case. In fact, the current projections by the current Social Security Trustees are that the fund will become insolvent in 2030, and that's causing people think about it.

The first approach that many people might take to this, if they have a totally clean slate, would be to approach something like risk theory where you have funds coming in and benefits going out. You might also have a ruin type of problem. That would be a very natural first step. But in fact the way people approach problems like this is to use statistical techniques, to use the past to say something about the future. You've got to have some kind of replication or something close in mind to that.

The first intellectual step that you have to take is not so much to take a look at the money going in and the money going out, but you have to break that down into components that you can forecast; in fact, contributions and benefit expenditures are things that are controllable by policymakers. It doesn't make a lot of sense to talk about forecasting those types of things that are controllable. What people will tend to do is break that down into pieces. There are two major components; demographic and economic. The demographics are things that actuaries deal with all the time—those would be mortality, marriage rates, and retirement. There are also a few things that actuaries tend not to deal with, but you would deal with them on a social insurance basis—for example, fertility rates and immigration. Oftentimes, it's the demographics that are the most important thing behind Social Security.

I'm going to cover the economic, in particular macroeconomic series. If you think about the way Social Security operates, there are natural increments every year in the amount of benefits that you receive that are indexed to inflation. We have this huge amount of money, around $440 billion; there are investment earnings associated with it. Then there is the wage index. Right now we're around $65,400. That's what we pay the Social Security contribution on. Usually it's split equally between employer and employee. The wage index is really important. Note also that you only pay contributions if you're working, so employment is another important series.

How are anticipated assumptions determined? They are determined not with a very formal mechanism, but based on negotiations among experts. These would be experts in Social Security and other experts in the Department of Treasury, Labor, and Health and Human Services. In particular, they're what we think of as deterministic; that is, they're not based on a stochastic model. Again, that's a very reasonable kind of thing to do. That gives you the basis for your best guess, for your estimate for the future. Naturally the kind of things that people like to know about are the sensitivity analysis—the what if. What if interest rates are a little bit higher or inflation is a little bit lower? Another thing that's published by the Social Security Administration are alternative projections below intermediate and high-cost assumptions.

That again provides you with that sensitivity analysis. That's the same kind of thing, however, you might get out of a probabilistic mechanism where we can provide things like forecast intervals, which provide you with information on how comfortable you should feel about things. This is just an alternative way of doing it, and it doesn't use the same kind of formal probabilistic assumptions. This is an alternative to that.

The thing that I want to show you is not about motivation; you'll hear more about the motivation later when you have the panelists talk about Social Security. I'd like to present a stochastic forecasting model of the economic assumptions, with the idea of applying this toward Social Security. Then I'd like to show you how there are certain patterns to the data. The patterns in the data will include auto-correlation. There will be some things that will not be based on the Course 120 and 121 syllabus: conditional heteroscedasticity, contemporaneous correlation and lead/lag relationships.

There's a complex model to do that, which you probably haven't seen before, but is something that is used a lot in econometrics. It's called a multivariate GARCH. It's a new model and it will be more fully described in the paper. In particular, I'd like to point out that once you have that model, and you have it calibrated, you can use it for forecasting funds. That will be the thing that you can apply not just to Social Security but to other types of things as well.

That was Part I, a brief glimpse at the motivation. Now the data. The data that we're going to be looking at are quarterly data beginning in 1953. Much of this is gathered just based on data availability. Good Consumer Price Index (CPI) data was available using quarterly as opposed to monthly because of discreetness in the data and a lot of other issues that were actually covered in a previous report by the Social Security Administration. We won't be talking about that quite as much.

I want to talk about how you validate these models. A standard technique that you might see in statistics would be to take a good chunk of the data, run numerous models, but then hold out a little piece of the data in order to see how the results actually work on real life experience. That's called an in-sample versus out-of-sample validation technique, and I'll show you that in just a little bit.

Just to give you a sense of this data, Table 1 shows four economic series. The mean inflation rate is 1.055% per quarter. The nominal investment returns then would be 0.5% above that. The wage rate would be around 0.25% in excess of the inflation rate. You can take a look at the summary of statistics, the high and the low, and then some standard deviations and maybe a mean versus median.

TABLE 1
SUMMARY STATISTICS FOR THE FOUR BASIC ECONOMIC SERIES
DATA ARE IN PERCENTAGES

Variable	T	Mean	Median	Standard Deviation	Minimum	Maximum
Inflation Rate	158	1.055	0.899	0.881	−0.553	3.927
Investment Returns (Nominal)	158	1.642	1.624	0.712	0.472	3.682
Wage Rate (Nominal)	158	1.312	1.333	0.731	−1.386	3.232
Unemployment	158	5.941	5.850	1.571	2.700	10.800

From the Floor: Inflation is 1% quarterly, but other data is annual, or are those all quarterly rates?

Mr. Frees: They are quarterly, with the exception of unemployment, which you would usually think of as being on an annualized basis; but then again, it's just the number of unemployed divided by the population.

From the Floor: So the inflation rate average is a little over 4% on an annual basis?

Mr. Frees: Yes. There are patterns in the data, and it's always difficult to communicate this to students when you do any kind of forecasting of the future. The game that you usually play when you're doing forecasts in a series is to identify the patterns, account for the patterns in the model, and then take a look at the things that are left over, which they call residuals. You then identify more patterns, and this can be called detrending of the series. It's always difficult to communicate that.

We start with a series, and we see patterns through the series (Chart 1). In particular, there are lots of patterns here with the investment returns—low, low and then the high stuff that we saw in the late 1970s and early 1980s. Unemployment and then wage rates look to be the most stable among the series. Again, why is it stable? A lack of patterns. It’s probably jiggling up and down more than anything else, but we call this the most stable in a sense that there are the fewest patterns to identify.

CHART 1
TIME PLOT SERIES OF FOUR ECONOMIC SERIES*

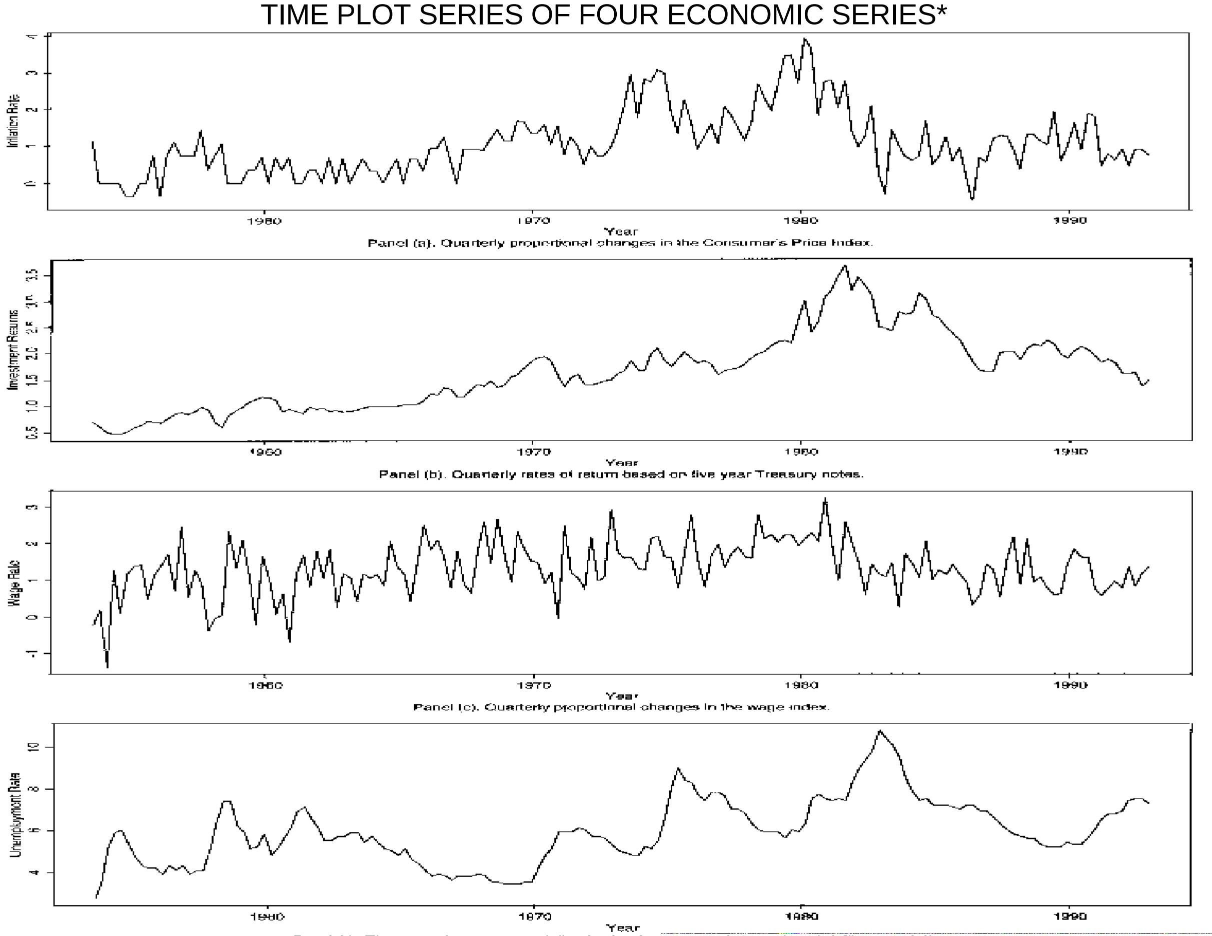

*Each series is over 158 quarters, from 1953:III to 1992:IV. Both Investment Returns and Wage Rate are in nominal terms, that is, unadjusted for inflation.

What are the kind of patterns that are here? For both the inflation rate and the wage rate there were the fewest patterns, so we could just use those series in and of themselves. However, for the investment returns and the unemployment, there were long periods where the rates stay the same and then suddenly jump. Let's just take a look at the unemployment series. There were long periods where the rate stays the same and then jumps up again. There are a number of different ways that you could explain that. You could either identify regimes or you could identify external forces to it. We want to talk about what are called naive models—things that you would use to make sense of the data in and of itself. That would serve as a benchmark for other types of things.

There are a lot of patterns here that you have seen in the data, and again there are a couple of different ways of doing this. One of the things that you would start to do or one of the things that we usually try to teach at the kind of Course 120 level would be to talk about the changes. One thing that works very well here in an actuarial context is to take a look at the changes or the differences of logarithms, because then once you go through a little tailor series of expansion you can argue that's very nearly the same thing as a proportional change. That's a very useful thing.

There are other kinds of patterns that you can look for. You can plot a series versus itself with a scatterplot. You can understand it using correlations. That's again the kind of thing that you do in forecasting. Table 2 shows the autocorrelations of the series or the correlation on itself. You'd do a scatterplot or again summarize that with the correlation. This basically says that if the inflation rate is high in one period, there's a strong tendency for it to be high the next period as well.

These would be the kind of things that students would see at a 120 level. The thing that we wouldn't do with 120 types of material would be to try to understand strong contemporaneous correlations (see Table 3). Contemporaneous would happen at the same time. Is the current wage rate affecting current unemployment? Is the current inflation rate affecting current investment? The answer is clearly yes. Do you go the other way and go to real rates? At the 120 level we'd like to somehow demonstrate strong intemperance relationships. Perhaps the best thing to do would be to focus on the correlations in Table 3. Chart 2 is called a scatterplot matrix. It's supposed to give you a way of seeing the relationships between two things. For example, the top scatterplot matrix would be a scatter plot of the investment return versus the CPI, a measure of inflation, and the correlation between the two is 19%, or 0.19. What's a little bit stronger is the relationship between inflation and the wages. It's around a 0.437.

TABLE 2
AUTOCORRELATIONS OF THE RETURNS: ORIGINAL AND TRANSFORMED VERSIONS

Lag	Investment Returns (Nominal)	Logarithmic Returns	Changes in Logarithmic Returns (y_2)
1	0.969	0.970	0.058
2	0.935	0.935	−0.135
3	0.906	0.906	0.073
4	0.873	0.874	0.086
5	0.836	0.837	−0.093
6	0.804	0.805	0.008
7	0.773	0.774	−0.166
8	0.750	0.751	−0.044
9	0.731	0.732	0.038
10	0.710	0.711	−0.053

TABLE 3
CONTEMPORANEOUS CORRELATIONS FOR THE FOUR ECONOMIC SERIES

	Inflation Rate (y_1)	Investment Returns (y_2)	Wage Rate (y_3)
Investment Returns (y_2)	0.192		
Wage Rate (y_3)	0.437	0.170	
Unemployment (y_4)	0.150	−0.368	−0.246

CHART 2
SCATTERPLOT MATRIX OF THE FOUR ECONOMIC SERIES

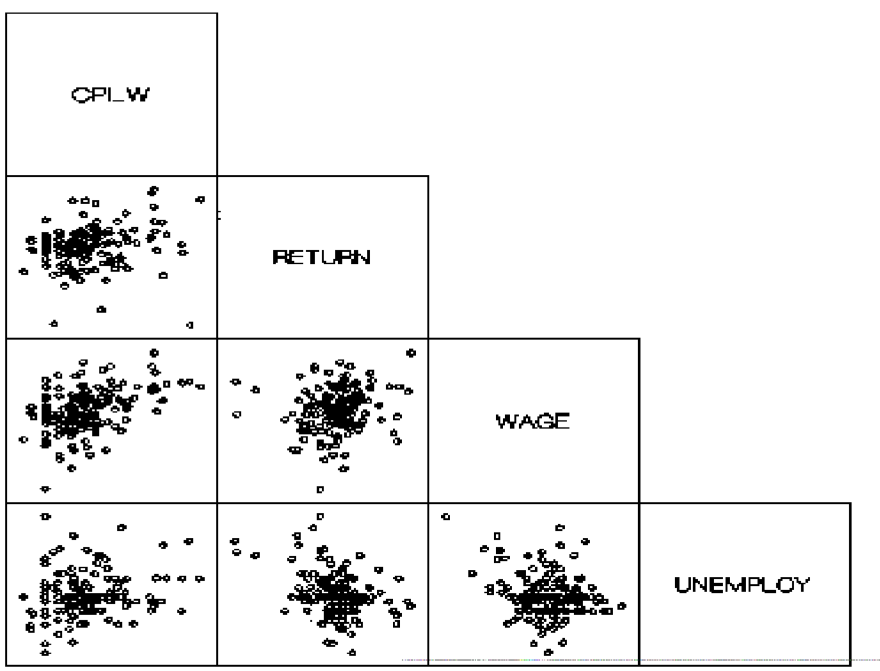

Looking back at Table 3 we see that inflation is positively correlated with wages. Inflation has less of a relationship with unemployment. The relationship between investment returns and unemployment seems to be negative for this data. In the paper we have the same kinds of relationships, but with real investment returns versus inflation, and that's a little bit harder to see and that doesn't take us down the right path, so I'm not going to show that to you.

One of the things we found out is that there are important lead/lag relationships in the data. You could also plot yesterday's investment return versus today's inflation. Yesterday's investment return starts to tell you something about inflation and that's one of the things that we get out of this. There are lead/lag relationships between the different series.

Then the last idea, which has been explored a lot in econometrics and studied a lot in the early 1980s, is the notion of conditional heteroscedasticity. That is basically a persistent volatility: you have periods that are relatively stable and then all of a sudden periods where there's a lot of movement, so large changes tend to be followed by large changes and small changes tend to be followed by small changes. That's the notion behind those kinds of models.

Those are the kinds of patterns then that you can see in the data. The next step would be to introduce some models to account for these patterns. Let me just remind you again why we do the modeling. We'd like to have a model that somehow represents the path so that we can project this model into the future. That's the whole name of the game. We're going to be looking at stochastic models so that we can automatically produce low, intermediate, and high forecasts. Then we can get a range of how reliable things are. We'll be looking at a time series model to explain things in the future. Many of these time series models simply take a current value in a series and take that as a function of previous values of a series. They are exogenous things—things that we can't really explain; they are errors or disturbances, or whatever you'd like to call them.

The next step then would be to allow the innovations; don't just have a constant variance, but allow those things to vary over time. One way to do that is with a conditional heteroscedastic model. That's something that you see in econometrics quite a bit.

One of the things that I did want to talk about is the standard stuff that we did. We identified the models to see which class that you have. Then we looked at various kinds of things. That's all done, we've written that portion up. I'd like to talk about how you use these models, because I think the inputs and the outputs are important. We also did a careful job in-between, but that's the kind of thing that you do in a quiet room.

You have these four economic assumptions. Having identified a model, you should use that for forecasting in the future. To help us choose models, we actually took a look at out-of-sample performance as well as in-sample performance. Again, we get a range of the variability.

The SOA is very interested in having us think about using economic models for a number of different problems. In a 1995 issue of *The Actuary* there is a publication of a survey that was done by an Economic Assumptions Guidance Task Force. I was on that task force. It was chaired by Godfrey Perrot. The basic results of that survey of the membership were that about 95% of the actuaries thought that economic assumptions are important for valuing our products and pricing them. Less than 2% of actuaries actually use economic assumptions.

Everybody agrees that it is important to take inflation into consideration. It's also important to take Treasury securities rates into consideration, but very few people actually use them. I'm going to be talking again about Social Security, but you could also use this kind of forecasting for a number of products. It would be up to you to determine how to use them price them the right way. That's the hard part.

Comparison of Stochastic Forecasts. We have two types of models. One would be just taking a look at each series in isolation of the others—that would be univariate. The other would be doing all four together at the same time. It's pretty clear when you take a look at the series that they do work together. However, a lot of times that you have real data and is it really worth the extra effort to bring all four things together at the same time.

When comparing stochastic forecasts, there are two types of models. One could take a look at each series in isolation of the others—that would be univariate. One could also take a look at all four together at the same time. It's pretty clear when you take a look at the series that they do work together. However, many times you have real data, and is it really worth the extra effort to bring all four things together at the same time.

If you're only trying to predict inflation rates, or if you're only trying to predict investment earnings, then you probably don't get that much extra mileage out of bringing the other three variables in.

Chart 3 shows two pictures. The top one is for multivariate forecasts; the bottom one is for univariate forecasts. Both plot the inflation rate; in both the actual data are the same. Out-of-sample validation is the same, but the forecast intervals are different; however, for qualitative purposes they're really essentially the same.

CHART 3
COMPARISON OF STOCHASTIC PREDICTION INTERVALS TO HELD-OUT ACTUAL EXPERIENCE AND TO SOCIAL SECURITY'S ASSUMPTIONS

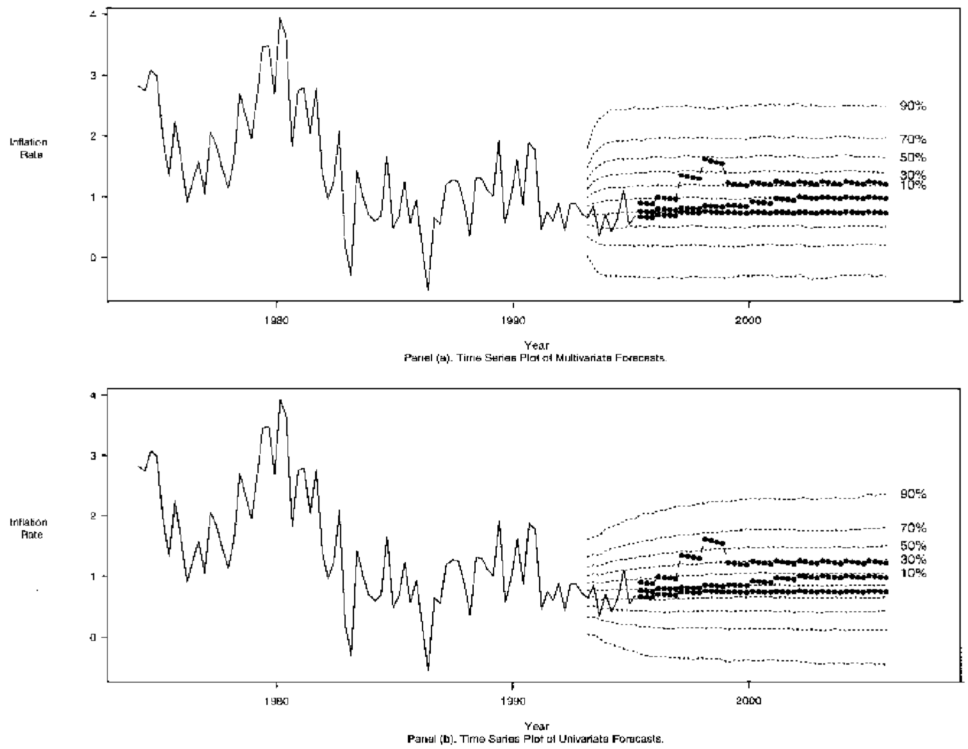

Panel (a). Time Series Plot of Multivariate Forecasts.

Panel (b). Time Series Plot of Univariate Forecasts.

The univariate forecast is a lot easier to do. You can look at inflation, in and of itself, and not even bring in other things. The multivariate forecast is probably more theoretically correct in the sense that you've accounted for more relationships, but again you don't get a lot of extra mileage. The only thing you try to do is forecast inflation or forecast any one of the four in the series.

It's very tricky, however, when you're trying to forecast all four things at the same time. To show an example of this type of forecast, we could create a little investment fund. I wanted to stay away from creating the actual Social Security Fund because we're still at an early stage for actual policy recommendations, but certainly you can talk about the notion of a Social Security Retirement Fund. This thing could be the Social Security system, but scaled down by a factor of 100,000. Instead of 42.9 million retirees, I would have 429 retirees with benefits. Instead of 139 million workers, we would have 1,390 workers. We would take that $436 billion and scale it down. It would be nice if we could do that to the real Social Security system.

The important thing, as far as forecasting, is we are forecasting something that's nonlinear in the sense that we're multiplying things in the future instead of adding things. We would take our fund, update it by one plus forecast interest, subtract that from our benefits, and multiply by the number of retirees. Then we would take the contributions times the potential workers and then subtract off those people who would not be available. It's a kind of recursive system where each component is being multiplied. It's a nonlinear system, but again if you are trying to write this stuff out, it would be a big pain in the neck, whereas using a simulation-based technique, it's much simpler.

From the Floor: Why would potential workers be used instead of actual workers?

Mr. Frees: Because the forecast of the population would be available. Again, you could say potential workers and then use that as a series and then forecast that directly. What we would like to do is have this fit into the demographic models where fertility, marriage, and remarriage rates can be forecast. Then they'll have very good estimates of the population, in the future. We're only interested in the working population so those forecasts would be available. We could just use our forecast of unemployment to get the forecast of the working population.

Chart 4 shows the very short, nine-quarter projection—comparing multivariate to univariate forecasts. The dotted lines are forecast intervals. This is for the fund. The bold plotting symbols are the actual experience. This is what really happened. We just held it out. We pretend it didn't happen. Again, this is just nine quarters so it doesn't prove anything. Maybe we just weren't lucky, but it does suggest anyway that the univariate forecasts are not doing as well as the multivariate forecasts. That's an important consideration in a sense that univariate forecasts are a lot easier to do than multivariate forecasts.

To do everything that we did with the ten-year projections, this is something that the Social Security Administration is required to do. We could have, but did not run the 75-year projections. It seemed to be kind of silly to run a 75-year prediction based on only 40 years of data. You could do that, but that's clearly where expert opinion becomes much more important. We weren't comfortable with that, but we were comfortable doing at least a ten-year projection.

The most important thing that I learned out of this was how tight the projections were compared to something like a 95% prediction interval. The future is probably a lot more variable than you might suggest by low, intermediate and high projections. Right now it looks like our forecast intervals, the point forecast, is higher than even the low cost, and the fund accumulates even more than any one of the three.

CHART 4
COMPARISON OF FUND FORECASTS
TO A FUND BASED ON HELD-OUT ACTUAL EXPERIENCE

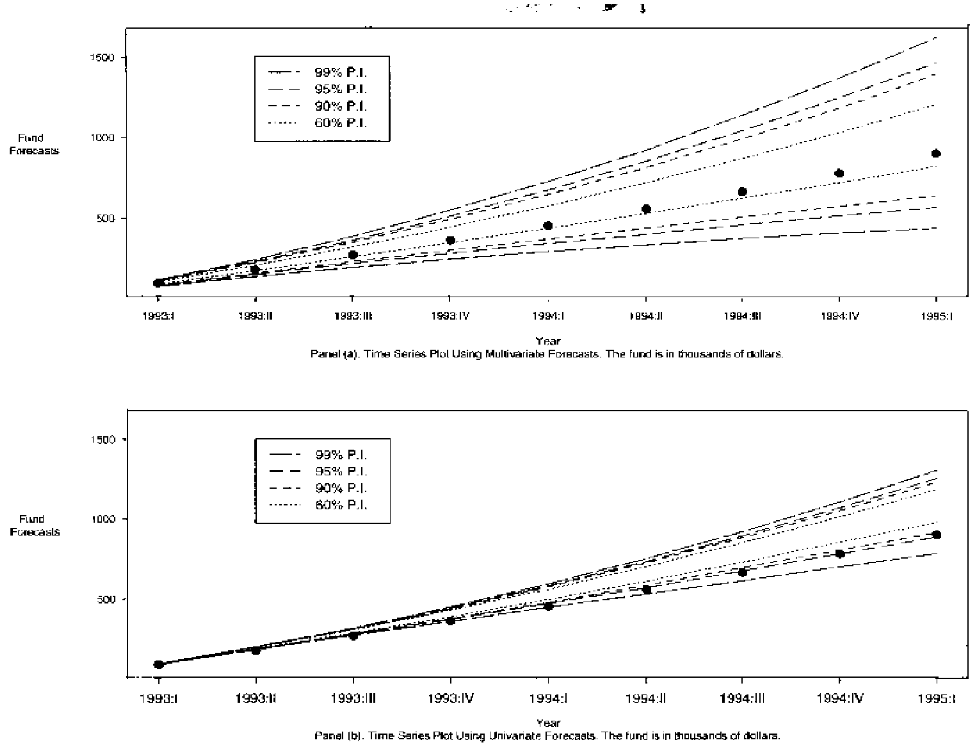

Panel (a). Time Series Plot Using Multivariate Forecasts. The fund is in thousands of dollars.

Panel (b). Time Series Plot Using Univariate Forecasts. The fund is in thousands of dollars.

This is the first stage in some other things that we are attempting to do. We have not built in the effect of regime shifts right now. Jenny and Margie, two of my co-workers, are working on that. The econometricians sometimes think about different regimes over which we have a low inflation and high inflation and how to model that. That's certainly something that we're putting in and there are a couple of other things with the number of parameters in this particular model.

Mr. Ostaszewski: Our next speaker will be Howard Young.

Mr. Howard Young: As Kris mentioned, I chaired the Technical Panel on Assumptions and Methods for the current Advisory Council. Every four years up until the current time under the Social Security Act, there has been an Advisory Council of citizens and others interested and knowledgeable about the Social Security System. They are not appointed because they are experts in the Social Security system, yet their role is to make policy recommendations to the government. They still have not issued their report. What they do traditionally is appoint a panel of experts—actuaries, demographers, and economists—who examine how the long-range projections are done under the Social Security system. So our charge was considerably broader than just the issue of this big stochastic projection question, but it fit into the overall situation. I'm going to speak to you specifically about the aspects of our report that related to stochastic projections. They were done without the benefit of Jed's work, but they did build very much on work that Rick Foster had done in the past.

All of the members of our panel essentially agreed on our recommendations, although Bob Myers, who was a member of the technical panel, did have some reservations about this particular aspect. If I don't cover them adequately, Bob, I'm sure you will. The other thing I can't resist commenting on is that Kris is one of the few people I could think of who, while we are in Disneyland, would remind us that we have to feel pain to understand things.

The other thing I should mention, triggered by Jed's comment about basing 75 years of projections on 40 years of data, is not having enough data. In fact, the big debate within our panel was not on whether we had enough data, but given all the data we had, which parts are most relevant? Over a 40-year time period, many things change and is it most relevant to look at the entire time period? Is it more relevant to look at the more recent period? We had enormous debates about that particular issue. On the one hand, that led us to some of these same points about the uncertainties of the projections and the need to look at new techniques, or at least to supplement what's being done with new techniques. We emphasized very strongly in our report that stochastic projections, while they are the more scientific techniques, they are no substitute for the need to reach some conclusions about what the basic assumption should be. I mean, you can look at the past data all you want, and you can analyze it all you want, and you can tease out all the correlations and all the variabilities that have happened in the past; nevertheless, you're stuck with the question of, how relevant is that for ten years into the future or for 20 years into the future or for 75 years into the future? While we believe and recommend that there are values to doing this, we emphasize very strongly that this does not really substitute or remove the uncertainty that one's left with when all is said and done.

A couple of other points. The concerns we had and one of the reasons that we recommended that stochastic techniques be used, not as a replacement for, but as a supplement to the current techniques, are somewhat related to the points that Jed made. The present method of giving an intermediate or best estimate of the future and a high- and a low-cost projection, does not give an indication, or at least an explicit indication, of what the range of variability is. Is the high-cost estimate equally likely with the low-cost estimate? What range of probabilities does that cover? And again while none of us really know, we can’t even take a guess at that.

A second point is that there is no indication of the relative significance of the different estimates. There is a range of estimates on mortality; there's a range of estimates on the economic variables; and there's a range of estimates on other variables; and there's no indication of how significant those are. For example, in one of the trustee's reports a year or two ago, there was a statement that the range of the mortality of variation had the most impact. That's true covering the range that they showed, but there was no indication about whether that range was more or less probable than the range that they had chosen to look at for the interest rates or the real age variables. We are trying to get a better handle on looking at those.

Also, there is no variability. The high-cost estimate says that everything, every single year in the future, is worse than was assumed in the intermediate cost estimate. The low-cost estimate says everything, every single year into the future, is better than for the intermediate estimate. There's no indication of which is more likely. Some of the variables might be better; some of the variables might be worse, but there's no mixture that one gets out of a kind of a stochastic examination.

We asked two members of our panel to put together an illustration to show (a) that it could be done and (b) what the illustration might look like if it were done—what might some stochastic projections look like if using the Social Security fund as a model. In effect, we jumped immediately to where Jed ended up in terms of turning out the long-range projections. We did look at it from the long-range point of view and not just from the first ten years. Although those particular models were published as appendices to our report, they were constructed by two of the individuals on our panel. They were, of course, discussed thoroughly with our entire panel, and as I say, most of us felt that it was worthwhile doing, and we explicitly suggested that this was simply an evolutionary way of moving along with the current techniques. We don't propose these as a substitute for the current techniques. We think that they have to be experimented with. We think that they have to be developed over the years. We think that they have to be compared and put out for discussion and evaluation by the professional community.

A very important point is, we have to learn how to communicate the results to the general audience. It obviously isn't enough to have this kind of a technical discussion of what goes on among people like ourselves who will follow it when all the policy decisions are going to be made by the policymakers, by the public press, by the media, by the electorate in general. We have to learn how to take this much greater degree of information that might come out of this and effectively communicate it. Those are some of the key points.

After the text section of our report is an extract from Appendix A, which is the model that was prepared by Martin Homer. Martin was a member of our panel.

He's an economist and he put together a model in which he focused mainly on the demographic aspects. He has since done some more work on the economic aspect and he looked particularly at fertility, immigration, and mortality. He used the intermediate basis that's done by the Social Security Administration as the mean of his model, and then he arbitrarily made the two sets of runs. In one he assumed that the range from the high-cost estimate to the low-cost estimate (that the Social Security Administration does) covers two standard deviations. On the other hand, he then made the again arbitrary assumption that it covers four standard deviations, so he characterizes these as the low and high uncertainty ranges of his model. Also, he made the point that because the Social Security Administration, which assumes that everything is either all good in the future or everything is not good in the future, implicitly assumes 100% correlation between the variables. He introduced some variation in the correlation. Those are really the distinctions in the various items that he came up with.

Table 4 and Table 5 come from the runs assuming the range from the high-cost to the low-cost estimate of the Social Security Administration covers two standard deviations or four standard deviations. In Table 4, perfect correlation is assumed, so you see that he has one for the correlation among these items.

TABLE 4
INPUT AND 2070 OUTPUT DISTRIBUTIONS FOR SSASIM RUN 21

	Fertility	**Immigration**	**Mortality Decline**
Mean Value	1.90	0.850	0.60
Std Deviation	0.15	0.100	0.30
Correlation:			
Fertility	1		
Immigration	1.0	1	
Mortality Decline	−1.0	−1.0	1

	Population	**Aged Dep Ratio**	**Total Dep Ratio**
Mean Value	368.0	40.8	83.3
Std. Deviation	38.0	6.3	2.3
Correlation:			
Population	1		
Aged Dep Ratio	−0.98	1	
Total Dep Ratio	−0.89	0.96	1

Table 5 shows zero correlation as distinct from the perfect correlation in Table 4 that we looked at. The wider range corresponds to run 21, or Table 4, that is, the perfect correlation, and the narrower chart corresponds to run 20, which is the zero correlation on Table 5. When there is zero correlation, of course, some things offset each other as time goes on because they don't all move in the same direction; therefore you end up with a narrower long-range projection than you do when there is with the same assumed degree of uncertainty, but with perfect correlation, everything is either going wrong or everything is going right.

TABLE 5
INPUT AND 2070 OUTPUT DISTRIBUTIONS FOR SSASIM RUN 20

	Fertility	**Immigration**	**Mortality Decline**
Mean Value	1.90	0.850	0.60
Std Deviation	0.15	0.100	0.30
Correlation:			
Fertility	1		
Immigration	0.0	1	
Mortality Decline	0.0	0.0	1

	Population	**Aged Dep Ratio**	**Total Dep Ratio**
Mean Value	368.0	40.5	83.0
Std Deviation	38.2	4.3	2.4
Correlation:			
Population	1		
Aged Dep Ratio	−0.69	1	1
Total Dep Ratio	0.36	0.39	

Let's turn to another model, which was done by Mike Sze, an actuary. Mike looked at a much broader range of variables. He looked at the inflation rate, the real wage increase, the real interest rate, the employment level, the labor force increase, mortality and fertility. Mike took the historical data and teased out of it the mean and the internal correlations as well as the variability for those. He used a number of projections, the size of the fund, the cost and so forth.

The trust fund ratio, for example, is under what he calls the baseline projection, which is using the intermediate projection. He took this intermediate projection as the assumed mean into the future and superimposed on that the correlations and the variability from the historical data (Chart 5). What you get is the trust fund ratio. The mean is the middle line, as you would expect. That's essentially what's being projected in the trustee's report. What he was trying to show was that the trust fund will run out of money around 2030. All that's really saying to us is that there is an equal probability under that assumption that it will run out of money before that date or after that date. We don't have any idea about the probability that it will run out of money within a certain range of years. Therefore, you could read this to say there's a 50% probability that it will run out of money somewhere between the year

2025 and the year 2035. That is information that you can get out of this analysis as distinct from the kind of analysis that is currently being done.

CHART 5
SOCIAL SECURITY BASELINE PROJECTION—TRUST FUND RATIO

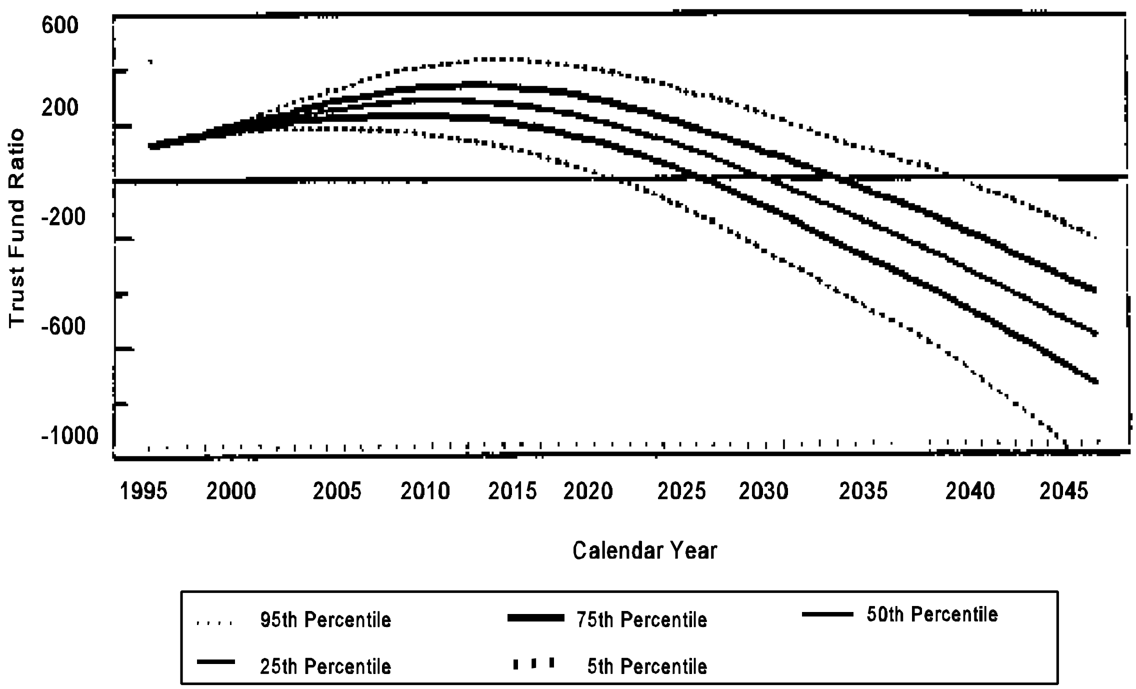

The final point I want to make is that when you look at any of these trend lines that come out of these stochastic projections, remember they are a summary of a large number of runs—in Mike's case, 1,000 runs. They are not a path that's being projected as an actual time path under any one situation. Under any one situation we wouldn't expect a smooth flow, but we would expect a lot of jiggling up and down, which could make a considerable amount of difference in how one views the stability of the results or the extent to which one might use them for policy purposes. With that I will try to answer questions later.

Mr. Ostaszewski: Our next speaker is Richard Foster, who is a source of some information that I'll try to verify right now. Did I understand correctly at the Bowles symposium that the combined Social Security/Medicare system will have a negative cash flow in 1998?

Mr. Richard S. Foster: Yes, in 1998. My purpose is to make a few comments on Jed Frees' paper and it's my pleasure to say that Jed and his co-authors are to be congratulated for a fine study that's really quite good. By way of background, this sort of approach to looking at Social Security and Medicare assumptions got started back with the prior Advisory Council, which also had a Technical Panel of actuaries and economists. That one was chaired by Steve Kellison who, as many of you know, is now one of the public members of the Board of Trustees for Social Security and Medicare. That panel recommended evaluating the assumptions that the trustees use in the low-cost estimate and the high-cost estimate. That panel suggested we don't have a very good idea of what those estimates represent. We know one is more pessimistic and one is more optimistic, but how likely are they? I have no idea, so they recommended use of statistic time series methods as an additional tool in the hopes that might help tell us about the range of uncertainty. That led to my own very modest contribution in this area, which is this actuarial study looking at the four, key short-range economic assumptions affecting Social Security.

The work that Jed did with his colleagues really took the next step that I had wanted to take and was not able to for a variety of reasons, and they thought it through much more effectively than I ever could have. I looked at univariate models, these four factors, one at a time, in an effort to find out what we could learn about the trustees' assumptions based on these statistical models. Jed and company were able to put these together in a multivariate approach and, as we've seen, that can tell us quite a lot compared to the simpler approach that I took. One of the key reasons that I really liked Jed's study was that it basically didn't find anything terribly wrong with my study. That was important to me.

What's the single most important finding from Jed's study? In my mind, when you combine the information, and you look at these variables collectively rather than one at a time, the result is more uncertainty rather than perhaps less. This was a source of some discussion and debate before your study came out and the question is raised in here several times. Would these factors acting together, where maybe one goes up in a bad way at the same time another goes in a better direction—would they tend to offset each other with the result being less variation, more certainty, or would the opposite occur? Would they add to the uncertainty and produce more variation than they might one at a time? What Jed's results showed was that when you compare the appropriate multivariate model with the combination of the univariate models done in what we refer to as a poorboy's method of combining them, You get more variation out of the multivariate approach. You get more variation, less certainty than you would through the use of the univariate models, and that's a terrificly important finding. Some of what Howard's Advisory Council Technical Panel did also supports that, but the approach that Jed took was sort of the classical statistical approach using the data itself, and, in my mind, is more meaningful in that regard. That was a very useful finding and has contributed significantly to the debate.

Now, of course, I also have a duty to point out the largest limitation of the study. It's not a major limitation, but it's significant and it's something, in fact, they're working on right now. When you look at past data and you take a statistical time series approach and you can apply the classical Box-Jenkins' identification techniques and so forth, just like Jed and his colleagues have done, the result can be significantly influenced by a kind of weirdness in the data. Let me give you an example of that. This is one key example, but there are others.

Chart 6 shows the increase in real average wages. This is where you take the nominal wage increases, net out inflation, and get a real wage increase. This is quarterly data. It's not exactly the same data, of course, that Jed used, but it's close enough. Over the entire period, starting back in 1950, the average increase in the real average wage was about the same as what Jed mentioned. The average increase was about 0.25% or 0.3% roughly over that period. Notice something that doesn't stand out here without the dotted line that I've penciled in for your benefit. There's a level shift in real earnings starting in about 1973. Prior to that level shift during the 1950 and 1960s, when the economy was recovering strongly from the effects of World War II, the increase in the quarterly average real wage was about 0.6%. However, starting in about the early 1970s things changed—the baby boom entered the labor force, and women entered the paid labor force. We also had an oil crisis. A number of things happened that had an effect to lower wage increases, and the average since the early 1970s has been about 0% per quarter.

CHART 6
QUARTERLY INCREASE IN REAL AVERAGE WAGES 1950: I THROUGH 1993: IV

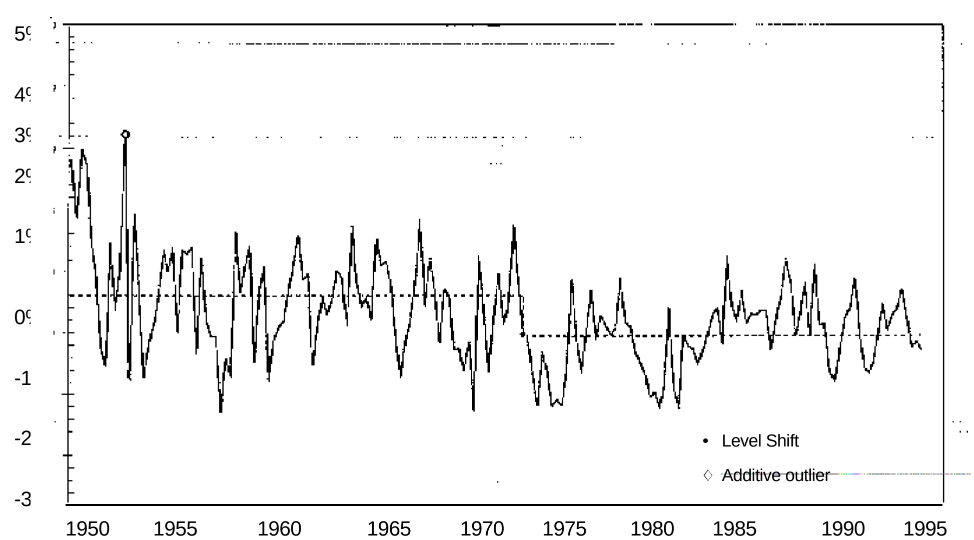

Now, that has implications, that we'll get to in a moment, for the nature of the model that you fit and the forecast that you make from that model. In this particular case, you can see that where we've been for the last 20 or 25 years tends not to be the same as the overall average of both of these two different periods, so that situation I would guess is the biggest limitation of what you've done so far. Jed characterized it as the regime shifts. It's the same issue. If you don't reflect these kinds of variations in the data, then you can get results that have no implications for your model and for your forecast.

The work that I did on a univariate basis did attempt to spot these things. They used a standard statistical test to look for level shifts and other outliers in the data. Then I attempted to build those into resulting models. I do not want to go through these things in any great detail. Listed below are the resulting models that I came up. The key things to notice are—in the unemployment rate model, the second one down, you see a model that shows the unemployment rate as a function of itself under earlier time periods. There is also a variable at the end there that stands for a level shift, naturally enough, and in the first quarter of 1974, the unemployment rate also experienced, in fact, a fairly dramatic level shift upwards in this case, where unemployment after that period has averaged about 7% and beforehand averaged only about 5%. That variable turns up with a very significant coefficient. The *t* statistic is almost three, so in fact using statistics, you can identify and test whether changes like this are significant or not. For the unemployment rate, it was. Similarly, for the real wage increase, where we showed you the data a moment ago, there is a level shift variable with a highly significant coefficient. It was a decrease with a negative sign, a decrease in the real wage growth.

Summary of fitted stochastic models:

Inflation rate:

$$c_t = 0.1342 + 0.5409c_{t-1} - 0.0010c_{t-2} + 0.3724c_{t-3} - 0.0474c_{t-4} + E_t$$
$$(1.78) \quad (10.69) \quad (-0.01) \quad (4.03) \quad (-0.67)$$

Unemployment rate:

$$u_t = 0.4749 + 1.3936u_{t-1} - 0.4931u_{t-2} + 0.2262LS_{74:\mathrm{I}} + E_t$$
$$(3.99) \quad (21.18) \quad (-7.60) \quad (2.92)$$

$$\text{where } LS_{74:I} = \begin{cases} 0 \text{ for } t < 1974\text{:I} \\ 1 \text{ for } t \geq 1974\text{:I} \end{cases}$$

Real Interest rate:

$$ri_t = 0.1731 + 1.0765ri_{t-1} - 0.1534ri_{t-4} + 2.972IO_{83:III} + E_t$$
$$(1.90) \quad (24.19) \quad (-3.44) \quad (3.83)$$

where $IO_{83:III} = 0$ for $t \neq$ 1983:III
1 for $t =$ 1983:II

Real wage increase:

$$rw_t = \underset{(4.10)}{0.3534} + \underset{(2.98)}{0.2138} rw_{t-1} + \underset{(2.00)}{0.1437} rw_{t-2} + \underset{(1.32)}{0.0954} rw_{t-3} - \underset{(-2.93)}{0.2017} rw_{t-4}$$
$$+ \underset{(3.61)}{2.3604} AO_{52:IV} - \underset{(-3.57)}{0.4011} LS_{73:II} + E_t$$

where $AO_{52:IV} = 0$ for $t \neq$ 1952:IV and $LS_{73:II} = 0$ for $t <$ 1973:II
1 for $t =$ 1952:IV 1 for $t \geq$ 1973:II

In both cases, the other two variables, the inflation and the real interest rate, actually had several level shifts, but the last one was the same as the overall average and so it will make no difference in the forecast, as we'll see. With the real interest rate, there were no level shifts identified in that model. Two very important ones, the unemployment rate and real wage growth, both had significant level shifts.

The handy way to summarize the results of the forecast you get from your statistical models is to make a series of forecasts and then take the five-year average of the 20 quarterly forecasted values and do that 5,000 or 10,000 times as part of a Monte Carlo simulation. You then end up with something like the distribution of the average rates of increase, like the average CPI increase or the average real wage increase. I got the four distributions of the average forecast and compared those to the Trustees' assumed averages over the same period.

For the inflation rate, the good thing was that the intermediate assumption used by the Board of Trustees was nearly identical to the median forecast for the model, so that was nice. The forecast range from the Trustees' Report captures what turned out to be about 50% of the variation that the model suggested would be likely to occur. A fairly similar result for the real interest rate, where again the Trustees' assumptions were fairly well centered relative to the model forecast.

Now the more worrisome thing turned out to be the other two that we already mentioned—the unemployment rate and the real wage increase. The intermediate of the Trustees' assumptions was well upward compared to the statistical model. The Trustees' most pessimistic alternative was close to the median forecast. These results, on a univariate basis, suggest that there is a strong possibility that, in the short range, the Trustees' assumptions in the two key areas are a little bit too optimistic, so that's worrisome.

How does that contrast with what Jed focused on? They found kind of the opposite. Remember, they found that the forecasts, based on the statistical model, tended to be better or more optimistic than what the Trustees were using, whereas what I just showed you tended to indicate the opposite. The difference lies in statistical methodology. If you have a model that does not reflect the level shifts or regime shifts, then, since these are autoregressive models and even in a multivariate basis, if they're sort of well behaved and I assume yours are, then the forecast from an autoregressive model eventually will go to the mean of the data. They are mean reverting models. If the mean of your data is based on the whole period, in this case it's about 0.3% per quarter, then your multivariate forecast for real wage increases will eventually go back to that overall 0.3%. If you build in the level shift in this case, then your forecast, wherever you start from, will eventually go to the latest mean, which in this case is 0%. Jed, if at your next stage, as you build in your Bayesian methodology and so forth, which I'm really looking forward to, if that ends up taking into account changes of this type, I suspect you will then end up with the opposite conclusion about who is more optimistic. Again my congratulations, and I look forward to the continuing work.

Mr. Ostaszewski: Let me say a couple of words and we'll ask you for comments and questions. What I found the most intriguing and valuable about the paper is the study of the internal relationships of all the variables, because I think that these kind of relationships are the key to what can go wrong and what we should worry about. I think that various internal feedbacks are the things that can go wrong. As we know, most countries in the world have some kind of social security system. When things go seriously wrong in them it's because of a drastic change in people's behavior in response to the features of the system. What we have seen in the U.S. over a probably similar period is a very big change in the savings behavior. We have disappearing savings. It's hard to say what caused that, but we definitely do have a problem. As a result of that, the distribution of the national income between labor and capital has changed. It used to be that about three quarters of national income went to labor and about a quarter went to capital. Correct me if I'm wrong, but I think it's something in the range of 65% to labor and 35% to capital. The share going to capital has increased dramatically and that's yet another picture of what we saw in the timid increases in wages. That's another face of the same coin really. This, I think, is something of the feedback that we receive from the economic system that we need to show in our models. That's why I think that this paper is very important, and we need to work on studying the relationships of various factors. Let me entertain your questions or comments. The first one is from Bob Myers, who needs no introduction.

Mr. Robert J. Myers: First of all, as I expected before I came here, I agree quite thoroughly with the comments that Rick Foster and Howard Young made. In addition, I would say that I think the stochastic analyses are very interesting, particularly to actuaries and even possibly the economists, whom I hold a grudge against. From a practical standpoint, my experience in dealing with policymakers, members of Congress and Congressional staffers, is that if you talk to them a lot about probabilities their eyes will glaze over and they say "Look, we've got to make a single policy; give us a single estimate." We always did give low-cost and high-cost estimates, so they at least were aware that some variation was possible, but from a practical standpoint, when you legislate you've got to view the single estimate and not pick some probability and use that estimate.

Second of all, when it comes to real wages, I think you should keep in mind that if the CPI is corrected technically (it has been a bit of an overstatement), the other side of the coin is good news. The real wages have really been higher than they've been shown to be.

Finally, there has been a lot of talk about the dependency ratio, and I think this isn't the world's greatest concept because it's taken on a static basis. I prefer the dynamic basis that the aged aren't people 65 and over; the aged are always people who are always two years older than I am.

Mr. Ostaszewski: Both of my parents are economists by education, and let me comment that there is really no need to hold a grudge against them because the third law of economic dynamics says that for every economist there is an equal and opposite economist. I know this from my family life.

Mr. Frees: Kris, could I just supplement Bob's comments about policymakers. In March I testified before a subcommittee of the Senate Finance Committee on our report and the chair of the committee was Senator Simpson, who is an ex-astronaut, and therefore clearly is fairly well educated. The thing he reacted to most dismissively was any discussion of stochastic analysis.

Mr. Robert L. Brown: Sometimes you get into theoretical discussions and a person will say, "That's very interesting." What that means is that you can probably publish a paper in a learned journal, but in the real world it will be deemed imperfect and even useless. Also, in presenting a mathematical argument about the future of Social Security, I get accused of demographic determinism by social scientists. That doesn't mean that I don't do stochastic models; it means that I haven't worked in behavioral responses, there are humans behind all of these models. If you were to have taken fertility rates from the mid-1960s to 1980 and project them, you would get to zero and somebody would sat that's probably not going to happen. I think you will find a new lower level of unemployment for the next 20 years because the baby boom is now in the labor force and, believe it or not, the increase of females in the labor force is going to have an upper bound. Again, if you were to take what happened between 1960 and 1980 and project it, you'd have 125% of all females in the labor force soon, so with all of that I'm wondering how you can build what a social scientist might say is a predictable behavioral response into what actuaries might like to keep pure and mathematical.

Mr. Ostaszewski: Rob Brown said that the behavioral response of people is something that can be—that is, very difficult to build into mathematical models. I would comment that there seems to be a lot of development in dynamical systems that is specifically interested in this kind of—let's put it, nonlinear changes; changes that don't continue along an established pattern. And there are areas in mathematics that specialize in this.

Mr. Young: There are a lot of people working in what is called adaptive behavior and some in a very complex way. These people try to look at things like how cells behave, but it seems to me you could have a model, in response to Rob's question, that builds in some feedback mechanisms and that recognizes that there are limits and things have to slow up as they approach those limits.

Mr. Foster: I'll comment further on Rob's point. It's not like any of us are proposing that stochastic models replace what's been done to date. What's been done to date has been based pretty much on judgment and informed judgment, taking into account exactly the kinds of things that you raised. These are not new issues. They've been talked about for 20 years. I see the stochastic models as an extra tool in the tool box. You can do all your judgmental evaluation; you can set your assumptions that way; and you can use statistics.

This may be a kind of a naive opinion, but if a model is way off base, then you worry about it. If it supports your judgment, so much the better. I think they can peacefully coexist.

From the Floor: One very simple comment. I am a little uncomfortable about how you use the word forecast. I think you should use the word projection because there's no policy change that you are evaluating here. You should consistently use the word projection given the current policy rather than forecast because you really don't know what's going to happen. A forecast is what actually happens, whereas a projection is your projection out given what the current law is. That's one point. Second, I have a comment regarding Bob Myers' comment about the CPI being overestimated, implying that real wages have been really high. Both reply to the wage series before and after the level change that Rick Foster depicted, so it's still true that the level shift in real wages has gone down.

Mr. Ostaszewski: By the way, there had to be some change if the distribution of income between capital and labor changed so dramatically, regardless of how inflation adjusted it. The total pie is split differently now between capital and labor than it used to be. You can see this in the stock market.