RECORD, Volume 24, No. 2*

Maui II Spring Meeting_
June 22–24, 1998_

Session 5PD
Designing a Stochastic Valuation/Forecast System

Track: Investment
Key words: Product Development

Moderator: RICHARD Q. WENDT
Panelists: CHRIS K. MADSEN
JOHN M. MULVEY†
Recorder: RICHARD Q. WENDT

Summary: This session focuses on a cascade structure model built for pension plans with a 15-year time horizon.

The capital market scenario generator is developed by:
- *Choosing an interest rate generator*
- *Building economic scenarios from the interest rate generator*
- *Building asset class returns*
- *Calibrating model to reasonable economic statistics*

Liability extensions are presented as:
- *Valuations across each modeled scenario*
- *"Stochastic actuary" development of assumptions across each scenario*
- *Minimization of model error estimates*
- *Asset/liability optimization*

Mr. Richard Q. Wendt: John Mulvey is professor of operations research at Princeton University and has been a wonderful asset in helping Towers Perrin in its systems. John will discuss Towers Perrin scenario generators and his view of them and how they're used in helping pension plans model their financial results. Chris Madsen will be talking about a brief overview of the forecast process with a couple of examples. I'll have some wrap-up comments on issues about setting assumptions for asset/liability forecasts.

†Mr. Mulvey, not a member of the sponsoring organizations, is Professor of Operations Research at Princeton University, Princeton, NJ.

Note: The charts referred to in the text can be found at the end of the manuscript.

Mr. John M. Mulvey: I will describe in a very general way the details of the Towers Perrin scenario generator, and also discuss how these techniques are used in planning. There are, of course, numerous applications of these techniques, and the question which comes up is, why do we need such complexity? Why do we have to develop a method that possesses multiple scenarios within a stochastic model? There are two answers to this question. First of all, the world has become market driven. The amount of money involved with the pension plans is tremendous, and investment decisions have a major impact on an organization's bottom line. We don't have the ability to smooth performance over time—if we make a mistake this year, we cannot correct for it next year. Secondly, new instruments are being developed, and these instruments have complicated cashflows. They possess implicit options and very complicated payoff strategies. To evaluate these techniques you will need stochastic methods. Chris Madsen will talk further about the creation of new securities and customized securities for pension plans.

Of course there are many applications of optimization throughout a financial organization. At the bottom of the risk ladder we see pricing models; these techniques help a company pinpoint its current surplus. If instruments are actively traded, pricing is easy. However, in many cases, the instruments are thinly traded, and pricing methods derive the value of these securities. Working up the organization you will want to include more details of the financial organization. At the top of the risk ladder are integrated risk management systems, which are becoming popular in financial organizations—banks, insurance companies, and other financial intermediaries. One step below are the dynamic asset and liability systems that come out of stochastic scenario generators. Note that information must be sent across different investment systems.

Here's a little bit of an ad. The described techniques are part of a Cambridge University Press book that should be out in fall 1998. This book describes the presentations at a conference at the Isaac Newton Institute and has additional examples. Three examples in the pension plan area are mentioned in the book *Towers Perrin System*. There's a stochastic system developed by Ortec that is used widely in the Netherlands. Frank Russell has a system which builds on stochastic programming as an approach.

Another good reference is by John W. Campbell, Andrew W. Lo, and Archie C. MacKinley *on The Econometrics of Financial Markets,* Princeton University Press, 1997. This book is very useful when calibrating these systems.

There are three components of an integrated financial planning system. First, stochastic models describe the uncertainties for the economic environment or for liabilities, your company structures which include calibrating and sampling. This component generates the scenarios for the simulation. The stochastic models start with a planning horizon. We project the economic factors across the discrete time horizon, and across a group of scenarios. The scenarios represent the future, and the analysis helps you understand the risk of various investment strategies. To get the right scenarios requires a calibration process and a sampling process. We must be careful about how the parameters are set in the stochastic model and also which scenarios are chosen to represent the future. These are critical aspects.

The second element of this integrated system involves the simulation of the pension plan over the planning horizon. Once the scenarios depict the uncertainties, you can simulate how the plan will react depending on the economic environment. The model should build the dynamics into the investment strategy to protect the pension plan's surplus. The strategies should respond to economic environments. The simulation requires decision rules that link investment strategies in conjunction with contribution strategies in order to manage the surplus. Once the simulation is put together, then optimization can be accomplished.

For scenario generators, we're basically setting up a system that starts with key economic factors, such as interest rates, inflation, and currencies. Bond returns are derived from interest rate changes. You must also link in other factors such as stock returns, how your liabilities react to inflation.

A modified timeline can illustrate the problem that as you move across time, the number of scenarios can grow rapidly as the number of periods increase. Interest rates, for instance, would have a branching structure, and for the scenarios to have a reasonable number of time periods, you would need to have a large number of scenarios, millions or hundreds of millions, if you were trying to be complete. The scenario depicts a representative set of the economic factors—the returns, the liabilities—across the time horizon, from today out to the end of the time horizon. We're not trying to predict any single path. Rather, we're trying to represent a variety of paths to determine what the errors and the dynamics might be as you move across time and across the economic environment.

We start with a fixed set of 500 scenarios. There is nothing magical about 500; one could tailor the number to the precision of the estimates that you would like to achieve. Then we have some time periods that go across time: *t* equals 1-*t*. Typically, for a multiyear model, the updates occur annually, and as you move further into the future, you might allow multiyear steps. From there, you take the asset returns, the liabilities, and contributions from these same factors. A key issue is to ensure that you derive the factors first and the returns secondly. If you go in the reverse direction—derive the returns from a multinormal distribution, for example, and look at the underlying interest rates that would have caused those returns, you are unlikely to have consistency.

The following equations are called the Orstein-Uhlenbeck process, and it's basically a model that indicates the change in interest rates, d_r, at time t, which is then a function of the current environment which is r_t, some parameter which measures the mean reversion level which is r_μ. The level that the interest rates will gravitate towards a_r which is the drift parameter and an instantaneous volatility parameter. This is a standard model for interest rates. It is one part of the Brennan-Schwartz model, and it's one of the key components of the Towers Perrin system. One could build other equations to represent the future, but this model is understandable and yet comprehensive enough to derive the range of interest rates as you would want.

Interest Rate Model

Short Interest Rates:

$$dr_t = a_r\,(r_\mu - r_t)\,dt + r_t\,\sigma_r\,dZ_r$$

Long Interest Rates:

$$dl_t = al\,(l_\mu - l_t)\,dt + l_t\,\sigma_l\,dZ_l$$

Mr. Mulvey: The dZ term is a random number between zero and one normally distributed. This term will allow the system to take stochastic derivations across time. As we move across time, the model may choose one path that will be high or low, coming out of the white noise term. The approach is called a stochastic difference equation. There are two equations, one for short, one for long rates, and the white noise terms are correlated. You'll get movement that move these interest rates across time, and you want to correlate the corresponding bond return.

Once you develop a model for interest rates and inflation, such as what we saw previously, then you can generate other parts of the scenario generating system. In other words, the model cascades down from interest rates and inflation down to the other elements. There is inadequate time to delve into the details. Suffice it to say that once you start with key factors, everything is consistently linked to those factors, whether its contribution policy or asset returns. Other issues can be included. For example, you can link the company's contribution plan based on the particular scenarios that you have in the system—inflation, wage inflation, stock dividend growth.

Scenarios that come out of stochastic generators ought to be realistic. Some systems that are arbitrage-free, for example, give results that may not be consistent and understandable. Also, patterns should be reasonable. As inflation is going up, interest rates should also increase. At some point, interest rates spike and then push down inflation. Understandability comes out of observing particular scenarios. Of course many things can occur, leading to the 500 scenarios.

Taking the 500 scenarios, we look at the distribution that occurs for cash and stock, which possesses a much wider distribution. In this case we measure the 90% level down to the 10% level. With 500 scenarios, we can't measure the tail of the distribution very well. So the range for stock for one year starts coming down. This is the compound return. Out here at 15 years, the stock market becomes much smaller in terms of its range and the cash begins to accumulate. Why is that? Why is the equity return having a funnel shape as time goes on?

Stock markets go up for a while. On average they come back down. That's part of the answer. The model possesses mean reversion in the stock returns, whereas if the model was a random walk, you'd see wider distribution as time goes on. You must decide whether the model has mean reversion or not. Certainly interest rates possess mean reversion, but cash is getting more and more volatile.

Mr. Mulvey: The cash series has a flexible mean reversion of a slight nature. It happens because of the inflationary increase, depending on the volatility of inflation and the link between interest rates and inflation. If inflation is volatile, then the cash return will also be volatile. If that cash will remain low over the next 15 years, inflation will likely stay low, then you have a higher band than you would show here. This is part of the assumption-setting that's built into these systems.

From the Floor: Do they all use the same parameters over time?

Mr. Mulvey: Parameters are periodically evaluated. There are three levels to evaluate scenario generators. One can evaluate the model structure, the equations of the model. Second, you can look at the calibration of the parameters. Third, you can evaluate the sampling errors and use standard methods for reducing the variance of the estimates. There are no real physical laws governing these systems. We can gain confidence in the methods at three levels: sampling error, calibration parameter error, and model structure error. We can evaluate and measure the sampling error and the calibration error. The structural error, however, is part of the stress testing. Chris Madsen will say a few words about this topic.

Calibration depends upon the target country. We compare two countries, the United States and Germany. In Germany the real rate of interest over the last 30–40 years has never dropped below zero, except for short periods of time, whereas in the United States there are negative real rates of interest, and, as a consequence, inflation goes up quite a lot. Looking at these two countries you have to decide to calibrate them differently if you believe that history is a guide to the future. Part of the challenge of employing these techniques is setting assumptions by looking backward to the past but also thinking about the future, as to what direction each of these countries is going in. That process is important when calibrating stochastic systems across different countries. There are different ways central banks operate, causing the calibrations to be different. I call Germany a Type I country. It's the only one that has this pattern. The United States is a Type II country. It has allowed interest rates to reach 9% negative real rates of interest. Other countries such as Turkey and other developing countries have allowed much lower negative real rates of interest.

Regarding the simulation, it should be robust over the set of scenarios—as an example, at each time period. The simulation makes a decision that will change the investment amounts. The time periods depend on the type of analysis. So a trader may look out using J.P. Morgan's system, maybe 10 days or two weeks, four weeks at the most, in which case the time steps might be half days or even hours. For a pension plan, you include longer horizons that update at each quarter or each month. There is no way to deal with the day-to-day volatility in the strategic system, so you're going to look at the time periods being fairly separate time, and, as I said, the scenarios depend on the situation you are working with. For 90% confidence, 500 scenarios seems reasonable; whereas, if you look at the rare amounts (10%), you're going to need more scenarios. There are modeling scenario systems for catastrophic events that are available that use 10,000.

We assume that decisions are not possible between periods. You make your decision and you basically live with that decision until the next point. We're just going to make our updates every time period. The decisions then are to try to make investment strategies and our liability strategies depending on the time period, the scenario, the types of asset categories, and if you have one contribution or you have one liability, but if you had multiple lines of insurance or if you have other kinds of liabilities, you have multiple decisions on the liability side. We must coordinate assets and liabilities.

There are numerous decision rules that one could apply. A traditional rule is fixed-mixed asset allocation, which is part of the Towers Perrin Oplink system. The portfolio returns to its target proportions every period. Let's say you have a 60%/40% rule. Every period you would update the model at the discreet points of time so that you start off every period at 60% stock, 40% bond. This decision rule is fine for long-term investors who have conservative strategies. There are dynamic approaches that look at surplus optimization or portfolio insurance with respect to surplus. These approaches are more dynamic than fixed mix and give different shape kinds of results. For example, you may wish to cut off the tail of the distribution on the downside—dynamic methods allow you to reduce the tail of the distribution. Once you select a decision rule, you simulate the rule and evaluate the result based on your investment strategy and the range of scenarios.

There are three basic approaches for dynamic financial optimization. There are no clear winners with respect to this choice of optimizing over stochastic scenarios. The three approaches are listed here. If you were an economist or electronautical engineer, you'd call this a dynamic stochastic control problem, and you'd use the equations of motion to develop a solution strategy. Brennan, Schwartz, and Lognado have an article in the Ziemba-Mulvey book, *World Wide Asset & Liability Modeling*, (Publications of the Newton Institute: Vol. 10), Cambridge University Press, 1999. Fairly straightforward problems fit into the structure. As soon as you get large complicated constraints, it's not going to fit into dynamic stochastic control. Only three or four state variables are feasible. In our case there are numerous drivers—interest rates, inflation, and other economic factors.

A second approach is called multistage stochastic programming. The scenario timeline branches out so that we can make a decision at every point in the timeline. This approach gives a much larger problem but provides much flexibility with regard to the conditional decision strategy. There are some difficulties because it requires massive computing power, and it is difficult to derive confidence estimates.

The third approach assumes decision rules and to optimize over those decision rules. This approach is more intuitive perhaps to investors. We saw fixed-mixed asset allocation. We mentioned portfolio insurance. People understand these rules. You can get confidence estimates. However, we get more complexity when we optimize over multiple time periods. Now, what does nonconvexity mean? I gave this talk a couple weeks ago in Greenwich at a reinsurance meeting, and somebody knew what nonconvexity was. They said, "Well, it's not convex." Non-convexity leads to a situation where the problem no longer has one peak. You have different islands with different peaks on it.

Traditional optimization methods typically will terminate at the top of a peak. For example, Excel states that optimization conditions are satisfied. What it means is it's at the top of a peak. It doesn't look out. You must employ nonconvex methods for solving these problems. You've got to have specialized nonconvex solvers. Even simple decision rules such as fixed-mixed asset allocation with expected value optimization would be nonconvex problems.

There's no clear winner at this point with regard to dynamic optimization. This area of research will lead to better dynamic investment strategies.

In conclusion, there are securities available that give very different behavior from the traditional stocks, bonds, and cash. One example is from Merrill Lynch that sells on the New York Stock Exchange or American Exchange. The security provides a guaranteed principal, a certain amount guaranteed every year, 1%, or maybe a fixed amount of money at the end of some time period. It also gives an increase of an index such as the Standard & Poor's (S&P) 500. You can readily buy its hybrid securities, and you gain a portion of the upside of the S&P, while limiting your downside to a fixed amount. You must have a clear idea of your goals in order to evaluate whether these securities are worthwhile to add to your portfolio.

There are several challenges. Ultimately, I think a growth area is customizing securities for a particular pension plan, individual insurance company, an individual, or buyer of a financial instrument. If you own an asset/liability system, you can identify bad scenarios, and reverse engineer which scenarios hurt you. Ultimately you can purchase instruments that will protect you against bad scenarios. This is surplus optimization, creating new instruments based on those bad scenarios. The second research area involves computing to get confidence in these scenario generators. How do we learn to trust these methods? Our analysis depends on these scenario generators. We can stress test them. We can look at historical data. We can calibrate them. But ultimately we have to find ways to develop confidence in these approaches.

Stochastic pension planning systems are now available. It is a matter of generating adequate support to put it all together.

Mr. Chris K. Madsen: I've worked with both John and Dick for a number of years, and I'm going to try and take you through some specific implementation issues. I'll try and build on some of the things that John covered in greater detail. I really find this to be a wonderful area to be in. I would strongly encourage any of you who have any interest to explore this area further. We really are just scratching the surface here. Once we get into more company-wide models, that's where the real value of this will surface.

There are basically two inputs to our models. We have the calibration parameters that come right off the stochastic difference equations that John talked about, and there are the economic assumptions. Economic assumptions start in today's environment. We refer to that as initial conditions. Then we define the normative environment. The normative environment is essentially the mean reversion point. I think someone pointed that out earlier on. Our model does have mean reversion.

Once we have all these scenarios developed, we can then build valuations, actuarial valuations, across each of these scenarios and get distributions of pension plan results. You use a similar process to get distributions of virtually anything that relates to capital market scenarios.

As John mentioned, we can optimize any reward or risk, and increasingly we've seen customized risk measures being in demand. People are no longer just looking at standard deviation of returns or surplus. They're really beginning to focus on downside risk or even much more customized measures. We had a situation—this was actually an insurance client—where it wanted to optimize volatility around zero surplus. It truly wanted zero surplus—which was an interesting thing. It's something we've never seen before. While that situation was truly unique, these customized risk measures are something that we see more and more of.

I'd like to explain the process that we go through. We work with four interlinking models. John mentioned the capital market scenario generator (CAP: Link) which lies at the bottom of all our models. Then we have the liability projection. Liability projection basically does a baseline forecast, and it also generates the sensitivity of your liabilities to salary scales, interest rates, and things of that sort. The capital market assumptions and scenario results are tied together by the financial forecasts (FIN: Link). Then we can take the capital market financial result, the financial results of the pension plan, insurance or anything else, and tie that together in OPT: Link. In OPT: Link we get an efficient frontier. We take asset mixes or strategies which will be a high-level view, but in most cases we work with fixed-mixed asset allocation. We then generate the distributions of the specific pension plan financials. We can get all kinds of output out of that.

As part of the process overview, John has covered the stochastic difference equations and the calibration parameters that we work with. I should point out that we work with a global model. We have ties across economic regions, and we work with a total of 804 calibration parameters for any one model. We'll generally set those once a year and then we'll only change the economic conditions, but our distributions will shift as the economic conditions change.

I want to highlight that we start everything in initial conditions, which makes sense because you want to make sure that we can tie into the pension plan valuation that was done previously. The stochastic difference equations then take us across each scenario, and when we get a little further out there'll be a bias towards normative conditions. There's no guarantee that we'll completely mean revert. That depends entirely on how the model is calibrated. The further away you are from your normative conditions, the longer it will take to completely mean revert. Every actuary has to make decisions every year in terms of discount rate, salary scale assumptions, and the works. We have to capture that process somehow as we move across each scenario. To accomplish this we designed a stochastic actuary that reflects your decision-making process. Of course we must specify contribution policy and a host of other things, including the employee population.

I want to talk a little bit about a couple of key issues, and John mentioned that multiple scenarios are valid because we get a range of results, the associated probabilities. It also makes it a lot easier to point to value at risk (VAR) which is becoming a popular measure, though generally it's more short-term oriented, but, with the distribution, VAR will normally use the fifth percentile, but that number will be readily available as will every other percentile. The error terms are normally distributed, though we make certain adjustments to them to give them the properties that we want. The overall distribution, therefore, is not normal. Stock returns are not normally distributed, neither are yield distributions. The easiest example is to look at yields where you clearly have a downside of zero, and there technically is no upside to the yield distribution. Further, we don't have an independent assumption from period to period which allows us to capture the type of distributions that John showed you earlier where you have the compound return of stock returns, the standard deviation of compound stock returns decreasing over a 10-year period, whereas you see cash showing expanding standard deviation over time.

I want to take a look at a specific example. In Chart 1 and Chart 2 I've picked a cash balance plan, stable demographics, projected unit credit funding method, and an ABO funded ratio with a little bit of a cushion. I did the same for current liability-funded ratio. In this particular example the company didn't make any contributions for 1998. I'd like to review the results using a 60/40 mix, just sort of the old rule-of-thumb, and then we can take a look at the contribution strategy and a few funded ratios.

In Chart 1, I've specified the minimum contribution being paid every year. That's, in other words, the plan sponsor policy. These are floating bars representing the 10th, 25th, 50th, 75th, and 90th percentiles. You have the mean indicated on top. What you see is they start at 105%, and then there's quite a drop. My first reaction was: how am I going to explain a drop of 15% in the funded ratio? I think this highlights one of the values that a system like this can bring to the table. In this particular instance it's because we're phasing out 1994. Nineteen ninety-four has much higher interest rates, and interest rates today are certainly quite a bit lower than they were in 1994, and so far they've been dropping. Even if we stayed at 5.9%, we're below that today; I think the long-bond yield is now 5.7% or something. Even if we stayed at 5.9% on average for all of 1998, we would see a drop of 47 basis points in the current liability rate, assuming they stay at the top of the corridor at 105%. For this particular plan that accounts for about a 10% drop in the current liability funded ratio. If the company changed its contribution policy from minimum to maximum, it would actually make a contribution in 1998, so they cushion the decrease a little bit and then, of course, it bounces back up. Other things are happening as we move further out.

But let's focus on the minimum contribution policy which tends to be the common choice, and I'm going to take you through an asset-only frontier and then an asset/liability frontier to capture some of the characteristics that are truly important to plan sponsors. Chart 3 is an example of the asset-only efficient frontier. Most of you have probably seen something like this in the past. This just basically demonstrates how we came up with the standard 60/40 that a lot of people point to. Sixty/forty looks fine in this context: 60% domestic large cap equity (S&P 500), and 40% bond index (Lehman Brothers government corporate bond index). Nominal annual returns are on the *y* axis, and standard deviation of those returns are along the *x* axis.

Now, if we take that into a surplus asset/liability efficient frontier, the efficient portfolios change dramatically. You will notice as I go through this that I've picked various risk measures. Chart 4 illustrates the versatility that we have. In this case the risk is shortfall, under 110%. Basically I've said that any funded ratio under 110% is risk to the plan sponsor in this case. Now, all of the sudden, the most conservative portfolio in asset-only space is the most risky, and not only is it the most risky, it also has the worst reward. Sixty/forty doesn't look so favorable anymore relative to the surplus efficient frontier.

From the Floor: Could you explain what the numbers 1 through 10 represent?

Mr. Madsen: These are different portfolios along the efficient frontier. The numbers 1 though 10 each represent an efficient portfolio for a given level of risk.

From the Floor: How many variables does that consist of or is it general asset classes?

Mr. Madsen: I think Table 1 can help me explain. On the top you have the asset/liability frontier, and on the bottom you have the asset-only frontier. On the asset-only frontier the most conservative bonds and cash, and I've actually restricted cash in this case because if I didn't, the most conservative would be 100% in cash, and I didn't see any pension plan realistically doing that. I have a couple of other limits in here. On the surplus frontier we see a heavy equity exposure. The most conservative here is 80% equity, but we are not recommending that plan sponsors go 80% in equity and suggesting that's necessarily a conservative assumption. The fact of the matter is, most plan sponsors don't have a long-term time horizon because they care about year-to-year fluctuations and the contributions—or charges to company equity that they're faced with in any one year.

I'll just highlight a couple of the reasons. If we switch to a shorter time horizon, even on the surplus frontier, we'll see more fixed income exposure. Plan sponsors tend to keep equity exposure lower because they're truly concerned with these bad events. What happens if we create a "bad" environment as opposed to the basic expectations environment? To analyze this I created the type environment that John was talking about—which is everybody's nightmare or the plan sponsor's nightmare—which is the Japan scenario where you have drastically falling interest rates and lousy equity returns. Japan is still down over 60% since its peak in the mid to late 1980s, and bond yields in Japan are down to one and a half percent on the long end currently, and they were at four and a half percent in 1993. If we have any chance of seeing that type of environment in the United States, then clearly that's something that plan sponsors should be concerned with, and we think they are concerned with it—even if it's a low probability event.

TABLE 1
ALLOCATIONS - DOWNSIDE RISK VERSUS ASSET ONLY

Allocations—Downside Risk Versus Asset Only										
Asset Mix %	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Cash-USA	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
30 YrBd-US	21.5	20.1	18.4	16.0	15.0	13.5	12.3	10.0	10.0	0.0
Eqty-USA	41.6	42.7	44.4	44.0	42.9	43.8	43.7	42.8	36.2	25.0
10YrBd-Eur	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Eqty-Euro	0.7	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10YrBd-Jap	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Eqty-Japan	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10YrBd-UK	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Eqty-UK	8.7	9.5	10.4	11.1	11.3	11.9	12.0	12.2	17.7	40.0
US Bnd Idx	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10Yr TIP	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	10.0
US SmCap	27.6	27.5	26.8	29.0	30.9	30.9	32.0	35.0	36.2	25.0
EAFE	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
WrldBndXUS	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Reward	116.5	116.9	117.3	118.2	118.7	119.2	119.6	120.5	120.9	121.4
Risk	6.2	6.2	6.3	6.3	6.3	6.3	6.3	6.4	6.5	8.9
Asset Mix %	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Cash-USA	10.0	10.0	10.0	10.0	10.0	2.5	0.0	0.0	0.0	0.0
30 YrBd-US	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	10.0
Eqty-USA	1.0	5.1	5.6	11.7	15.5	17.7	20.7	27.1	30.1	45.0
10YrBd-Eur	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Eqty-Eur	0.0	0.0	0.0	0.0	1.2	1.7	2.5	4.6	5.6	0.0
10YrBd-Jap	21.5	14.2	7.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Eqty-Japan	0.0	0.0	0.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10YrBd-UK	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Eqty-UK	1.0	4.7	7.6	12.8	14.3	16.0	17.4	19.9	21.1	0.0
US Bnd Idx	60.0	60.0	60.0	53.8	43.5	44.5	38.6	21.2	13.1	0.0
10Yr TIP	6.5	4.5	3.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0
US SmCap	0.0	0.6	5.5	11.7	15.5	17.7	20.7	27.1	30.1	45.0
EAFE	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
WrldBndXUS	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Reward	88.2	91.9	95.5	102.1	105.3	108.2	111.2	117.0	119.8	120.6
Risk	26.2	22.8	19.6	14.3	12.3	10.6	9.3	7.9	7.5	6.8

For the stress scenario, again 500 simulations, we've shifted the equity return and bond yield distributions left. It's just a standard probability distribution. I should tell you the expected annual return of equity under basic expectations is roughly 10%. Under the stress scenario, the equity return is roughly negative 5%. In case anyone is looking for historical precedents for this type of environment, we saw this several times during the 1970s over a five-year period, and we saw this several times during the 1940s and, of course, during the 1930s. The average T-bond yield drops by just 50 basis points. The average T-bond yield under basic expectations is 5.9% while the average is 5.4% under stress.

We create an efficient frontier under stress. If I told you that the equity market was going to do terribly over the next number of years, and that bond yields will drop substantially, then, it's a no-brainer. Then everyone would be in fixed income, and you wouldn't mind having a couple of puts in your portfolio, and that's exactly what we see here. These frontiers over here are heavily dominated by fixed income, and some have a put option included. We see that the portfolios optimized under normal conditions fail. We think that any plan sponsor strategy needs to recognize the concern the plan sponsor has with this type of event occurring. In essence, even though there is a low probability of this event, that probability has extreme value to the plan sponsor, so it should be tested.

To address the concerns of plan sponsors, we need a strategy that works best under both stress and basic expectations. In this case I've included an at-the-money put on the S&P 500 Equity Index. The put option is turned over every year. I should point out that we've done this type of analysis using different types of options and derivatives, and the put is actually not efficient under basic expectations. Under the basic expectation scenarios, the put option is not an efficient asset class, but an at-the-money bond call is. But in the interest of keeping things simple, let's focus on the put option.

I picked an alternative (to 60/40) portfolio and then went back and took it through FIN: Link to get the pension plan statistics on this portfolio. It's very similar to the original 60/40, though you could argue it's more conservative. In this case we actually have 50% fixed income instruments, 40% long T-bonds, 10% in the new Treasury Inflation-Protected Securities, and then some international diversification and, of course, 3% in the put option. Basic characteristics here: longer duration, better diversified, and downside equity protection.

When we compare the alternative portfolio to 60/40 and look at the contributions under stress, we see that under the alternative portfolio the whole distribution shifts down, so the expected contributions drop by two million 10 years out. Similarly, we can look at the current liability-funded ratio under stress which, of course, is substantially better using the alternative portfolio because of the put and the fixed income exposure. We conclude that the alternative portfolio works under stress, but what happens under basic expectations? Ultimately we have to show that the final portfolio works statistically under basic expectations as well as under stress. Fundamentally, you'd expect there to be some cost of this strategy unless you find a great, new asset class like those that John alluded to: such as the new derivatives that are coming out from Wall Street firms such as the Merrill Lynch MITTS or

Solomon Brothers index-linked securities. Indeed, we see that the alternative portfolio return is lower than 60/40 under basic expectations, reflecting the costs of the past.

The compound return is where you see the true cost. When you have 3% of your portfolios in a put and the equity market keeps going up, it is clearly going to cost you something. We loose roughly 160 basis points annually for compound returns under basic expectations, but the point is that it all depends on your perspective. If your perspective truly is an asset/liability perspective, then you shouldn't really care about this asset only sacrifice!

It is prudent to review asset allocation, funding method, and all the things that the plan sponsor is concerned with under these what-if scenarios. We've just touched the surface, and there are many financial products that are being developed and can be developed that can help plan sponsors better protect surplus risk in the future.

What I'd like to discuss is setting assumptions for stochastic modeling, and there are three key issues I'd like to talk about. First one is looking at history, and many analysts like to go back to history, pull out the financial statistics, and the historical numbers as a starting point for setting their expected returns and volatility assumptions. I'm going to show an approach.

Second, coordinating international assumptions is a key issue. The Towers Perrin model is a multinational model, and we are able to look at currency from different points of view. We can create a model from the United States point of view, United States currency from the German point of view, or Japanese currency from either the United States point of view or the German point of view, all simultaneously. We're getting into that with respect to multinational companies. We're modeling all their subsidiaries in a coordinated way worldwide. There's a recent paper that gives some interesting thoughts on that and shows the interconnections among all the currencies.

Finally, a big issue is its linkage of bond returns to starting economic conditions. Many analysts tend to look at history and say average tutorial bond returns were 8% or 6% or whatever and use that as the assumptions for their capital market simulation. My point is that if you're starting a simulation at January 1, 1998, or any other time period, the key determinant of future bond returns will be your starting yield, and I'll show you some examples of why that is true.

Chart 5 is an example of how the analysis of historical data can change. It shows the same data in two different ways. Depending on how you look at them, you can get very different results. The thick line is a pattern of bond returns for calendar years starting in 1981. When we look at calendar years over that time period, the standard deviation is 14%. That's what you'll see in the standard books that report financial data. However, if you were to take that same data on a monthly basis and simply start one month later, (i.e., starting in February 1981), and slice everything up into 12-month increments, you would get quite a different picture. It happens that when you start in February, still looking at annual returns but not on a calendar year basis, the annual return now has a volatility of 12%, and this is only for a one-month difference in the starting point. There's really an important issue of not just looking at the calendar year but getting into the variability that shows up when you look at the details on a monthly basis. Chart 6 shows that in a lot more detail.

Chart 6 slices everything up to monthly starting points, but each dot represents annual returns, in a sense, of 12 monthly returns grouped into 15-year periods. It's the standard deviation of a 15-year period but simply starting at either January, February, March, etc., through December. Over those periods ending in the last 13 years the standard deviation has been as low as 10.50% and as high as 17%. There's a significant difference in the numbers depending on exactly what data you're looking at and how you slice them. Again, if you look at the traditional textbooks, such as Ibbotson, they show it on a calendar year basis, and they show very little of the variability of the statistics.

The point of this section is to say when you're setting your expectations for long government bonds returns, there may not be a single number that is the right number. You may have to look at a range of numbers, and any number within that range may be an appropriate number and justifiable.

The next segment pertains to currency modeling. It's based on a very interesting paper that was in the January/February 1998 issue of the *Financial Analysts Journal*. (The article was "Maintaining Consistent Global Asset Views (With a Little Help from Euclid)" by Singer, Terhaar, and Zerolis.) The authors showed a visual portrayal of the relationships between the currency of three countries—U.S., Germany, and Japan. The insight of the authors is that there is a trigonometric relationship governing standard deviation and correlation of currencies. They show a triangle with each of the three countries at a vertex. The length of the lines between vertices represents volatility of one country to another. The cosine of the angle at each vertex represents the correlation between the currencies of the two other countries.

For example, if you're looking from the U.S. perspective at Germany and Japan, then an increase in the correlation between the mark and the yen implies a reduction in the mark/yen volatility (i.e., the volatility of the yen from the German perspective or, equivalently, of the mark from the Japanese perspective).

This model shows that the assumptions for volatility and correlation must be consistent. If you change one element, you're implicitly, and perhaps unknowingly, changing some of the other elements. I urge you to look at that paper as it takes a complex subject and shows how the relationships can be explained in a simple way.

The final issue I'd like to talk about is linking bond return assumptions to starting economic conditions. The long bond yields of January 1998 were about 5.9%. If you look at history and assume that future bond returns over the next 10 years are going to be a compound return of 7.5%, it doesn't sound like that's too dramatic an assumption. Even though you're starting from a 5.9% current yield, 7.5% return is just a little bit above that. You might think it's a little optimistic, but it may be reasonable. In fact, it really takes some extreme changes in future bond yields to develop a 7.5% return, and that's because the yield change tends to fight against the price change to get to total return. For instance, if yields were to drop to 4% over the next five years and they stay for the next five years, you would get positive price changes for five years. On the other hand, your yields would be lower, and the tension between the two would give you a net result of 7.5% over a full ten-year review.

From the Floor: Could you explain that again?

Mr. Wendt: If over the next five years your yields start at 5.9% and drop to 4%, you're getting a very nice performance over those 5 years because you're getting price appreciation as the yields drop. At that point, though, your yield has leveled off at 4%, and your return for the following 5 years is basically 4% a year. So, the combination of the 4% for the last 5 years plus the positive price changes for the first 5 years gives you a net of 7.5% over the total 10-year period.

The other example is a little more extreme. If your yields increase for 5 years and then level off, they would have to go up to 29%. In other words, if yields increased to 29%, you would get losses for the first 5 years, but then you would have a level 29% return, approximately, for the next 5 years. The combination of those losses and the 29% return for 5 years will again give you 7.5%. The 5 years is arbitrary, and if you were to take a 10-year trend instead of the 5-year trend, it makes the upside a little more challenging. You would have to have a trend towards 53% yield over 10 years to have a net return of 7.5% over that 10-year period. So that's obviously a pretty implausible scenario. If the yields went down, they would only have to go down to 4.5% over 10 years to wind up with a 7.5% return. In fact, when you look at all these possible scenarios it turns out that the only plausible way to get a 7.5% return is either a temporary increase in yields followed by a drop back to where you started or a slight drop to start with and then staying level.

You really do have to pay a lot of attention to the initial yields. One particular reason for that is that if you're doing an asset/liability forecast, your liabilities are normally going to be tied to the initial yields, initial inflation, and the initial starting conditions, as are the assets in that pension fund. To somehow disregard the initial conditions and say, well, we think bond returns should be 7.5% or higher is really creating some dramatic assumptions beneath the scenes as to what you're really implying for future yield changes.

Looking at history in the right way, don't just take one number, but look at all the different relationships. Look at currency. Look at the linkages. Get that paper from the *Financial Analyst's Journal*. It was really an interesting paper. When you're doing the asset/liability forecasting the starting point is really critical in many of the issues.

From the Floor: John, what types of advanced optimization are you looking at now in your work?

Mr. Mulvey: The research has two elements. For nonconvex approaches, it's a matter of trading off the time to get global solutions versus the precision that you want in those procedures. From the stochastic programming approach, you will need much more efficient solvers. At this stage I think it's premature to apply these methods. We favor optimizing decision rules, and using nonconvex techniques to solve the resulting problem.

Mr. Josiah Lynch: John, your reference to bad scenarios under customizing products for individual pension plans raised a couple of questions. One is, are you talking primarily about reflecting the liquidity requirements of the plan or are you talking about long-term solvency?

Mr. Mulvey: It would depend on the situation you're working with. It could be neither case. If you were trying to minimize your contribution over time, then it would be an issue of trying to protect the downside on the plan surplus. In that case you could pick out scenarios that have large contributions over time and develop instruments that would protect against that.

Mr. Lynch: Going on to a question for Chris on the same topic, is it customary or rare to reflect the plan liquidity needs in the forecasting?

Mr. Madsen: No, I'd say that's fairly common. Clearly, if you have a plan that is scheduled to have a substantial amount of payouts for whatever reason, that needs to be reflected. For the distributions you saw we just showed you a sample, we could also generate those for payouts under the given asset allocation strategies. In your particular case the risk measure might be one that's related to the payout strategy.

Mr. Lynch: Chris, is it common or unusual to alter the package of turnover assumptions, salary scales, to match up the demographic implications of the investment assumptions on the individual plan valuation?

Mr. Madsen: Could you clarify what you mean by the demographic—

Mr. Lynch: Suppose you have a hundred scenarios. If you're talking about a high inflation rate in one of them, is that reflected in the salary scale simultaneously in doing the liability forecasting?

Mr. Madsen: Yes.

Mr. Lynch: OK. Is turnover altered according to the economic assumptions?

Mr. Madsen: The turnover of demographics as opposed to the asset turnover. You're talking about demographic turnover. In other words, employee population.

Mr. Lynch: Right. If you have high inflation, you typically have high turnover because people are job-hopping.

Mr. Madsen: Excellent question. No, we don't do that.

Mr. Wendt: I can add to that answer a little bit. The Pension Benefit Guarantee Corporation actually is developing a very complex economic model of large pension plans and plans that it's insuring under its pension guarantee insurance. In its model it is going into great detail with the assumptions and having corporate bankruptcies modeled in accordance with the capital market assumptions and simulation results. The demographic assumptions would change in certain cases with the economic assumptions, as well as the size of the employee population. That's one of the more advanced models, and there have been a number of actuarial discussions on that. There's some literature on that.

Mr. Lynch: I was curious to know if this was sort of automatic or whether it was special customization in consulting.

Mr. Wendt: At this point it's pretty much nonstandard, I would say.

Mr. Andrew Cairns: I've got a comment about some remarks that John made about which method or optimization you should choose. You had a few possibilities. Some of the work that I've done is based on the dynamic stochastic control. It offers a nice complementary service to the other methods of optimization. What you mentioned was that it tends to be based on relatively simple setups because dynamic stochastic control is a very complex problem that requires small models that you can cope with in terms of computing power. But the benefits that you get out of doing that before you perhaps go to some of the other methods is that you can find out from these methods precise answers in terms of what is optimal for a very simple model. What you would then hope is that when you move to a more sort of complex world, more complex liabilities, a bigger set of assets, more complex asset dynamics, you hope that what you have found to be optimal in the dynamic stochastic control model will also be nearly optimal in the more complex world when you apply one of the other two optimization methods.

Mr. Mulvey: I think that's right—stochastic control methods can inspire decision rules. For instance, the fixed-mix rule is optimal under certain conditions. The problem is that empirical methods are limited. If you use a structure for solving general asset/liability models such as the Brennan-Schwartz-Lagnado, you still have to solve the final system of equations and have a large finite element or some dynamic programming. Certainly for inspiring decision rules it makes good sense, and I think you'll find hybrids of these three methods in the future.

CHART 1
PENSION PLAN FINANCIAL PROJECTIONS
XYZ, INC.
1998–2007 CURRENT LIABILITY FUNDED RATION, MINIMUM CONTRIBUTION POLICY

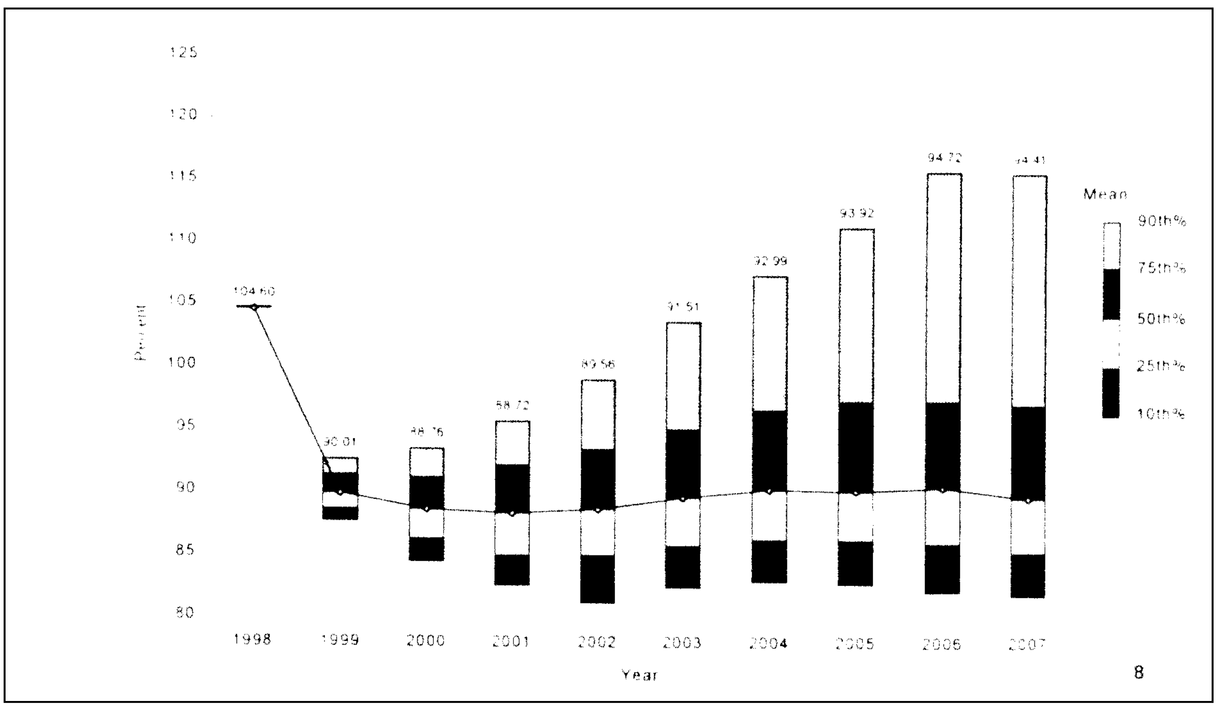

CHART 2
PENSION PLAN FINANCIAL PROJECTIONS
XYZ, INC.
1998–2007 CURRENT LIABILITY FUNDED RATION
MAXIMUM CONTRIBUTION POLICY

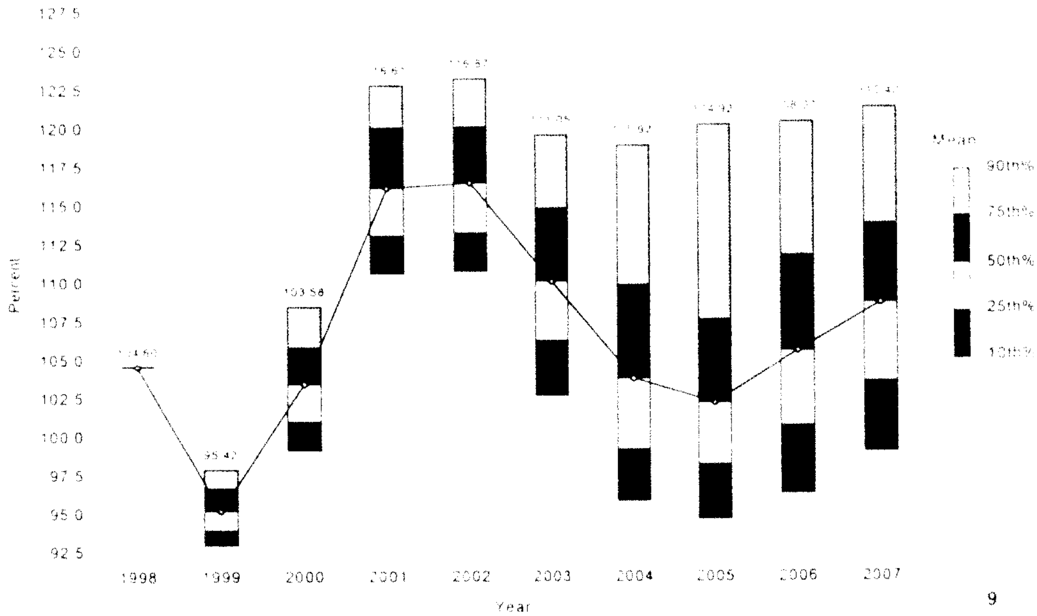

CHART 3
TRADITIONAL ASSET ONLY
ASSET-ONLY EFFICIENT FRONTIER
10-YEAR TIME HORIZON

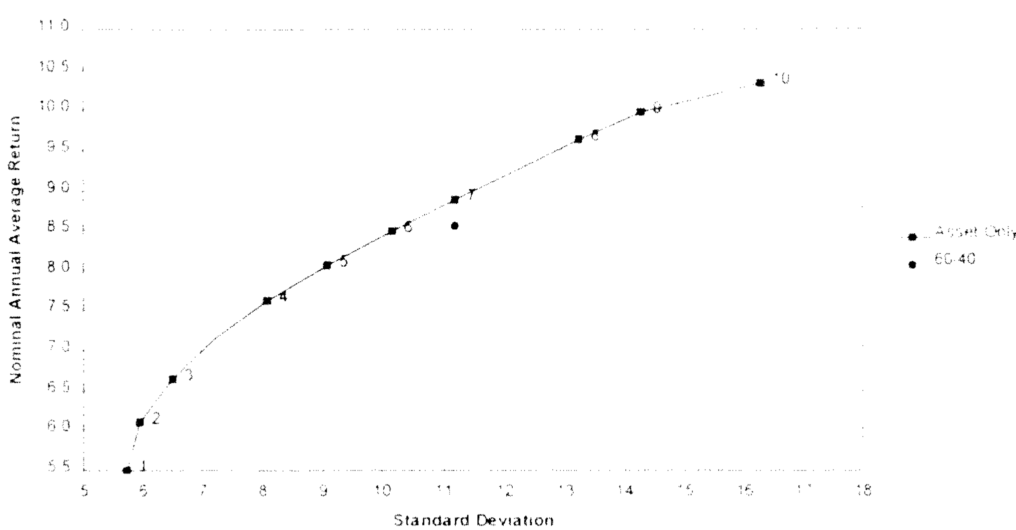

CHART 4
ASSET/LIABILITY FRONTIER USING FUNDED RATIO DOWNSIDE AS RISK
ABO SURPLUS VERSUS ASSET-ONLY EFFICIENT FRONTIER
10-YEAR TIME HORIZON

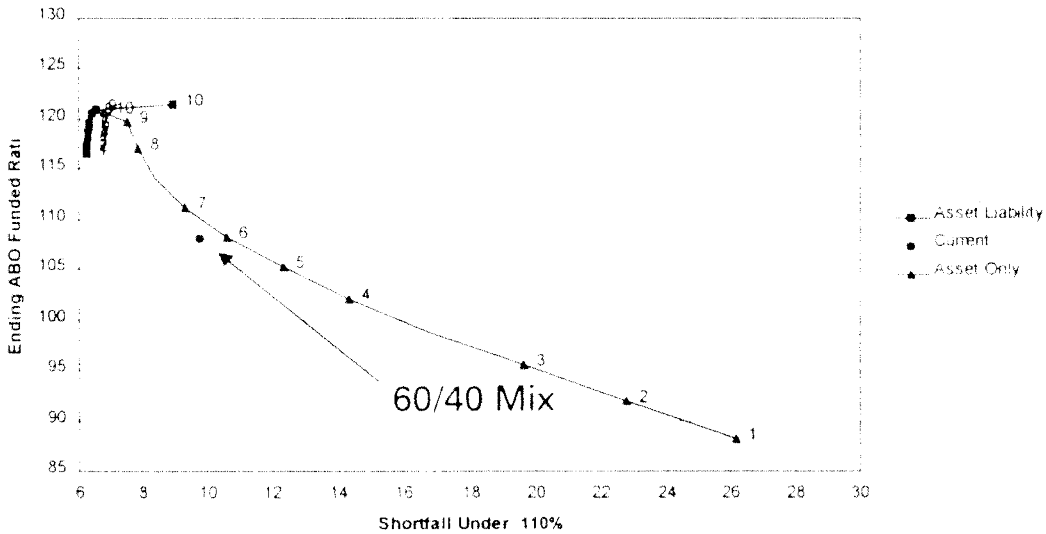

CHART 5
LOOKING AT HISTORY—A SHIFT OF ONE
MONTH CHANGES THE STD. DEV.—14% VERSUS 12%
COMPARISON OF HISTORICAL RETURNS
LONG GOVERNMENT BONDS

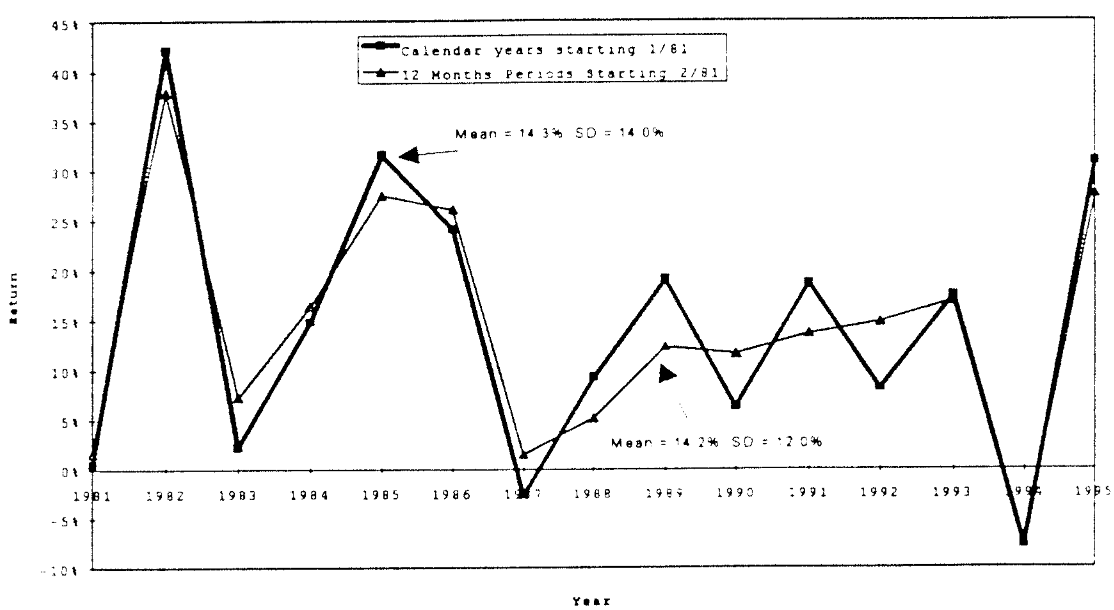

CHART 6
LOOKING AT HISTORY: STATISTICAL MEASURES
CAN VARY SIGNIFICANTLY

STANDARD DEVIATION OF LONG GOVT. BONDS
OVER ROLLING 15-YEAR PERIODS