**1998 VALUATION ACTUARY
SYMPOSIUM PROCEEDINGS**

SESSION 36PD

REAL-TIME STOCHASTIC ANALYSIS

Frederick W. Jackson, Moderator

Mel Stein

Larry White

REAL-TIME STOCHASTIC ANALYSIS

MR. FREDERICK W. JACKSON: This session was planned in conjunction with another session on new frontiers and asset/liability modeling (ALM). What we have is a little different. Our two main panelists are really going to talk about the perspectives they have on real-time stochastic analysis.

We want to discuss how asset/liability management is synonymous with a dynamic financial analysis (DFA) and quantitative risk management. It's an evolving art and science, and these folks on the panel are going to talk about a good part of the science. They have some system solutions that are innovative and entrepreneurial, and they're going to talk about those. *The Dynamic Financial Condition Analysis Handbook* was really the first substantive ALM literature that tried to pull things together. It was intended to reflect contemporary knowledge and practice at any given time. That's being updated regularly, and I was involved with the original handbook. I've been recently involved with the ALM principles task force exposure draft that has just come out. It tries to focus on principles avoiding practice. It takes a step back. All the people on that committee are involved heavily in the practice, but wanted to step back and take a look at the principles.

There's a definition given in that exposure draft: "Asset/liability management is the ongoing process of formulating, implementing, monitoring, and revising strategies related to assets and liabilities in an attempt to achieve financial objectives for a given set of risk tolerances and constraints." Specific mention of practice is avoided. This session addresses practice. We're going to focus on innovative system enhancement approaches rather than the new conceptual tools that were discussed at the session on New Frontiers in ALM.

The panelist's focus will be on the systems implementations which include asset risk exposure optimization given a company-specific set of liabilities. I think Larry White will be talking about that from applied quantitative solutions:

Mel Stein will be focusing on total risk exposure optimization, which includes stochastic analysis of alternative asset and liability strategies. There's overlap between their two presentations, but I think you'll find them quite interesting.

In conclusion, I think you'll want to consider that the principles work emphasizes that there's no one right way to do ALM or DFA work. I think you'll want to be clear on the menu of evolving concepts and tools that you have to choose from. You want to start with a clear direction or get help constructing a company-specific risk management process before moving forward.

MR. LARRY WHITE: I'm with Applied Quantitative Solutions (AQS). We deal with asset/liability management, and I'll leave it at that for the moment. My partners and I, actually came from the asset side of the business. We started back in the dark ages of the 1980s, back when we used to look at yields in a book. We didn't even use a computer. Many people didn't have computers on their desk. I even had a lot of clients without fax machines, so the state of information technology has evolved greatly. I also think a topic like this might not have even been understood ten years ago, and today it's pretty heavily attended. I think that says something about where the state of this industry and this technology is going.

When I was speaking with Rick earlier in the year about this particular topic, he mentioned real-time stochastic analysis. I thought it was a mouthful. I started to break down that term. How can we apply this? I think we all know what analysis is. Unfortunately, *analysis*, by its very definition, suggests a very time-consuming process. By its very nature, it's almost contradictory to real-time. Now we hear *real-time* bantered about in terms of quotes by and large. You hear it on the Internet and in reference to stock trading type tools, if any of you have access to that.

We're going to talk about something that we have developed over the course of the last five or six years that is as close as we can come to real-time, given the limitations of computer power and of some of the database engines that we have to work with currently.

This process models balance sheet and income statement performance. Many people have asked me, "Do you optimize yield? Do you minimize covariances and all these very heady topics? When you get into a board meeting and you explain that to the board members, they don't know what that stuff is, what do they want? They want to make more money, but is maximizing yield going to make more money?" There's probably a good correlation, but they're not necessarily the same thing. Wall Street has known and exploited that for years. This optimizes specific financial statement performance benchmarks. As such, we can say things to our process like make more money, make more surplus, make more earnings. Because it has been reversed-engineered into balance sheet calculations, it does indeed accomplish that. It produces results in an asset management context.

The way that we are predominantly applying it is as a precision asset targeting vehicle. The state of the art today or up until now has been one where we do all this work on the liability side of the balance sheet and the income statement. Then we boil it all down to something like duration and convexity and we hand that to the asset managers. They can't do a lot with that, and do you know why? It's because every asset manager in the country gets 20 or 30 phone calls a day from every broker in the country that wants to sell them a security because that's how the brokers make money. I know I did that for a dozen years. The thing is, the human mind simply cannot process all this information at once, at least not the mathematical information or the quantitative data. A benchmark like duration and convexity and these sorts of things are valid, to some extent, and certainly that's the way most assets are managed today, but there is a better way and that's what I want to show you.

How does this process work? Liability financial information is extracted from client models. We currently have interfaced with *ALFA, TAS* and PTS, we extract relevant data from the income statement, the balance sheet and the cash-flow runs. We essentially model the income statement and the balance sheet within our own model. The assets are modeled discretely. We use Capital Management Sciences (CMS) bondage. That is not the "end-all" asset engine. There are others like Global Advanced Technology (GAT), *Barra*, Intex and Yield Book. My apologies to those I have left out in coming up with these rosters, but those are some that I'm sure you're all familiar with. The asset and liability data are pre-processed to verify the model accuracy. In other words, the internal AQS model is set up such that it looks almost identical to the outputs from TAS or PTS or

ALFA so that we know that our model and the model that you're used to seeing and, more importantly, that your board is used to seeing, are lined up. We know we're doing the same thing because if we optimize for something that is not your business line, then we've wasted a lot of time.

Next, constraints are specified and objectives are defined. We designed as many constraints as we could reasonably come up with. In the three dozen or so applications that we have accomplished, we have found that there is not a constraint that we cannot work into our modeling language. It's fairly robust.

Results presented are real assets in real-time, and that's important because we go to asset managers all the time. We have these asset sector allocations—25% mortgages, 50% bonds, and 30% corporates. That still leaves much room to foul things up. What's more, we've seen a lot of things that are in the form of proxy assets. Proxy assets are not really out there on the street; therefore, the asset managers can't really go out and buy them. So they have to do a great deal of estimation to see if they can hit those targets.

What we mean by real-time is that we will take dealer inventories on a daily basis, and we'll run those through CMS bondage or the client's software. We'll run each *CUSIP* for every scenario considered, and we generate data, book value, cash-flow, market value and accrued income. We do not calculate duration in the beginning, we do not calculate convexity, and we do not analyze the data. We simply generate the raw data.

Finally, we verify the results independently, and this was necessary in the beginning because when we developed this process, no one wanted to believe in the AQS model. They asked, "Who are you guys?" This is in large part due to the fact that there's already a CMS, *TAS*, PTS, and *ALFA* interface. Once we have an optimal solution set of assets, this time zero portfolio, we can regenerate those cash-flows for that optimal portfolio. We can zip those up, e-mail those to the client, the client puts those through *TAS* or PTS, and they can verify in their own modeling language that the results that we are showing them are indeed superior results.

The graphical case studies that follow represent results using current asset targeting versus optimal targeting. I broke the slides up into before and after. We will represent *statutory* surplus development. We've done economic surplus development. We represent distributable earnings, and if there's something else that your company is a benchmark for or that represents making more money or making less money, we can represent that to the solver. These graphs are actually prepared from data generated by *ALFA*, PTS and *TAS*. We did not use our data to generate these graphs; this is the stuff that is post-verification. It has gone through *TAS* or PTS, and we've extracted those data. They've been prepared, taking particular care to apply consistent assumptions and constraints. If I took a AAA portfolio and knocked it down to a BBB portfolio, that wouldn't be a very relevant comparison. I figured you were smarter than that, so we are keeping the overall portfolio waiting in exactly the same way. What's more important is we're keeping the quality distribution identical. If the portfolio only has 2% in *BAA-1*, we constrain the optimal portfolio to be *BAA-1* or only at a 2% maximum.

The asset targeting grid example is probably the heart of the entire system and that is something that the asset manager can look at and say, "This is as much information as can be put in my hand to communicate to me exactly what the objectives of management are and what the results of the actuarial evaluations are as well."

What do these results mean? You do quantitative asset targeting and you have improved risk management. You will see that shortly. Better risk management makes for improved free surplus and capital distributable earnings (DE). Better free surplus, capital and DE improves investor relations because you can credit more if you like. Better investor relations means you have more assets to target and you can grow your business.

Chart 1 is a representation of where we find ours fitting into the world. On top, we see the asset engines: GAT, *Barra*, Intex *BondEdge*, and *Yield Book*. On the liability side is PTS, *TAS* and *ALFA*. Each of these systems engines will make a claim. I know CMS does. They claim that they can optimize too. We have begun to move away from the term optimization. It's overused and everybody thinks they know exactly what you're talking about when indeed it can be something

CHART 1
The ALM Problem: Communication

ASSET / LIABILITY MODELS PROVIDE RESULTS BUT NOT SOLUTIONS

INTEX

BOND
EDGE

YIELD
BOOK

GATT/
BARRA

OTHER

MODELS

PORTFOLIO MANAGERS
CHIEF FINANCIAL OFFICERS
CHIEF INVESTMENT OFFICERS

ASSETS

AQS
ALM
SOFTWARE
SOLUTIONS

THE ALM SOLUTION

LIABILITIES

ANALYSTS
ACTUARIES
ALM DEPARTMENTS

LIABILITY MODELS

PTS

TAS

ALFA

DISCRETE ASSET MODEL

DISCRETE ASSET MODEL

DISCRETE ASSET MODEL
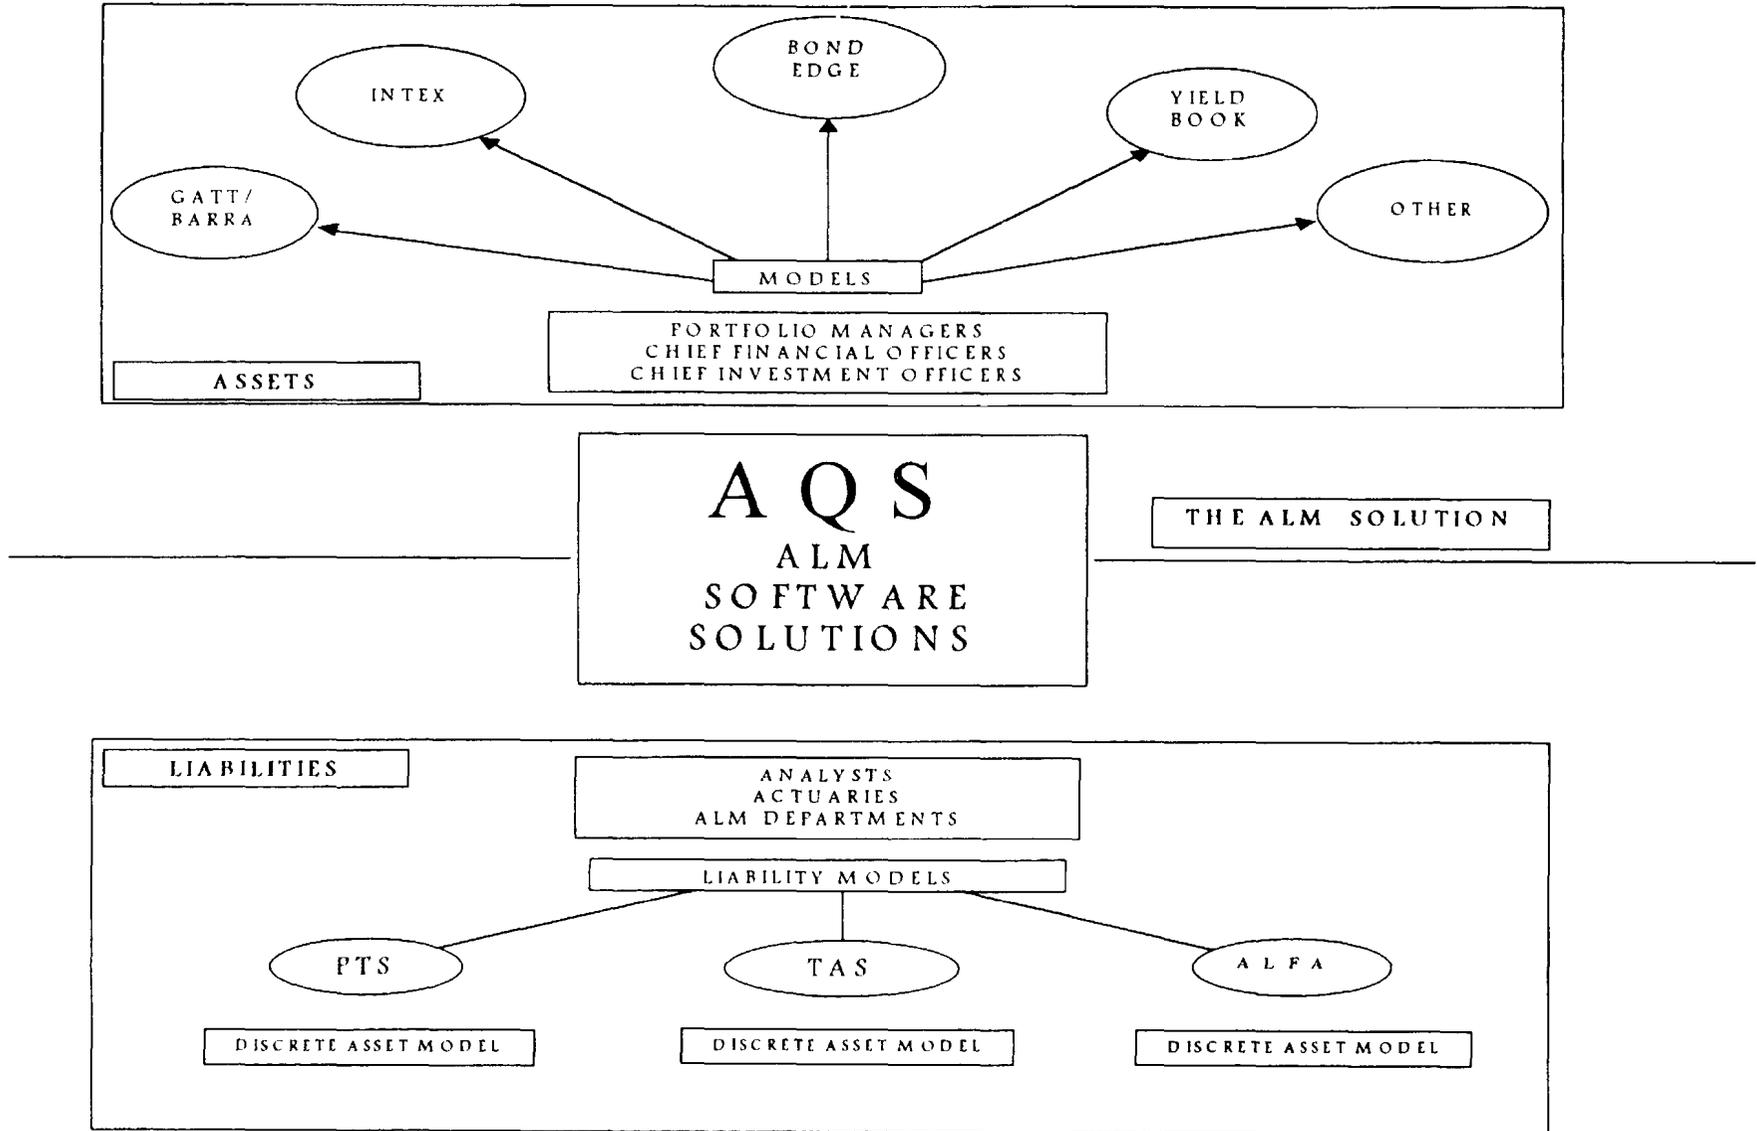
different to everyone. On the liability side, I know that *TAS* and PTS and *ALFA* will represent themselves as being asset/liability tools. Like Rick was saying early on, asset/liability management is somewhat ill defined so it sort of is what you make it.

We'll get into some results now. I'll give you a rough layout. Chart 2, current surplus development, shows a particular line of business. We have millions of dollars on the side so we have –$2 billion at the bottom, about $1.5 billion on the top. There are seven scenarios there. When you're looking at seven scenarios, that's likely the New York seven scenarios. We've done many more scenarios. We've done a number of stochastic scenarios. Our process is not dependent upon any particular group of scenarios. As each of these are color coded in most terms or it used to be such that in the upright scenarios, you had mostly the negative surplus development, the down-rate scenarios, you had the up surplus development. With the high guaranteed rates of some business lines now, they're seeing the reverse. We have very low interest rates, and so we hit a 5% guarantee on an annuity line. For example, we're starting to see that cut into operating profits. As rates go down, there's an even bigger problem than was ever dreamed of back when we had really high rates. In this particular case, this business line was going negative in the up-rate scenarios.

Chart 3 is the proposed surplus development. This is what happened after we ran the optimization and came up with an asset targeting. Again, these assets are very similar to the original portfolio and structure, distribution, position size, and virtually every item that an asset manager would look at. In fact, the asset manager plays a vital role in developing this process. The important thing to note in Chart 2 is the scenarios were divergent. In Chart 3 they tend to converge in the after case. Isn't that risk management? We don't want to make a call on where we think interest rates are going. Nobody is that good. There are a few that are that good, but if you are that good, you go out and trade your own account. You don't need to be working at a company. We are doing risk management here, and as such, we are trying to identify all of the possible scenarios and optimize against those.

CHART 2
Current Surplus Development
(TAS Output Treasury Reinvestment)

$MM vs YEAR

Legend: s1, s2, s3, s4, s5, s6, s7

Y-axis: 1500, 1000, 500, 0, -500, -1000, -1500, -2000

X-axis: Initial, 1997, 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017
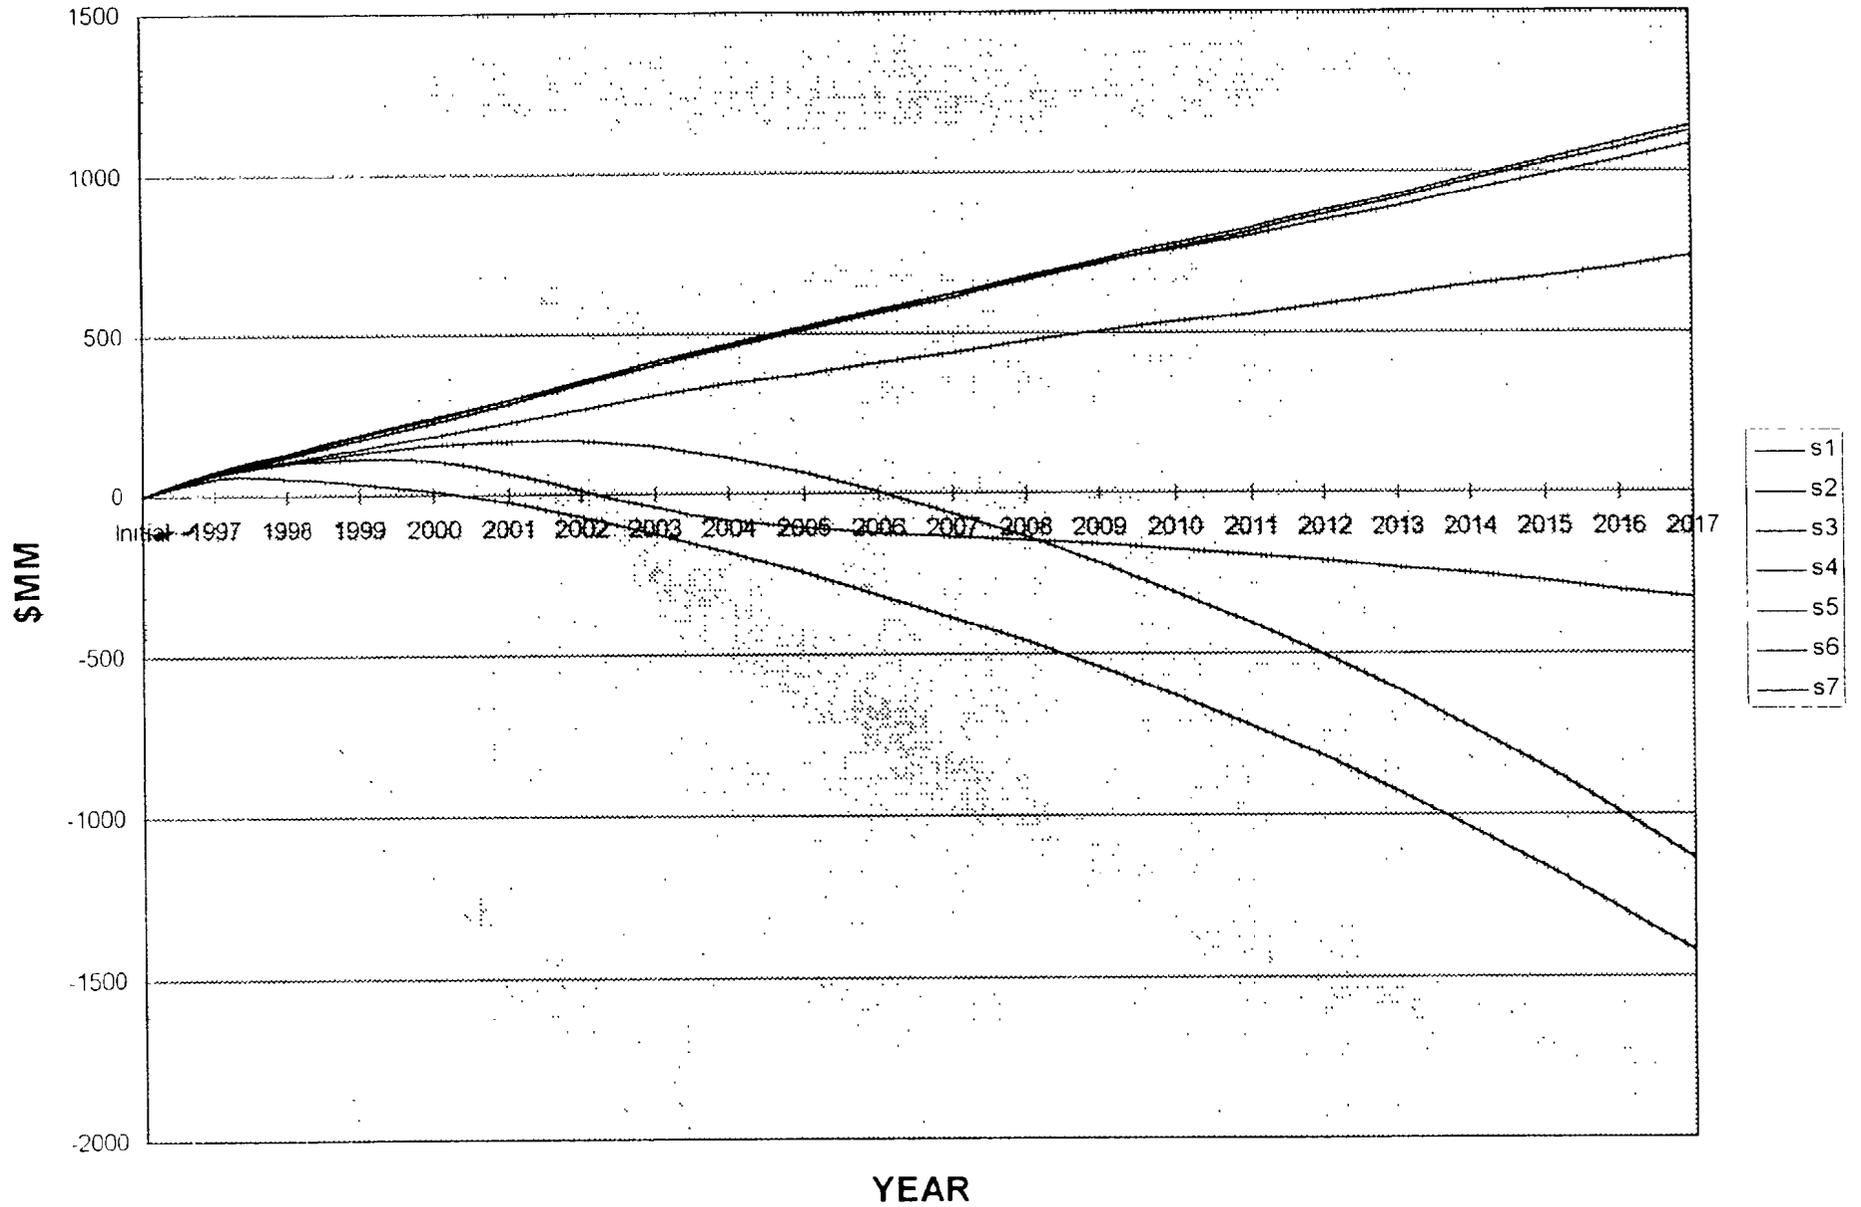
CHART 3
Proposed Surplus Development
(TAS Output Treasury Reinvestment)
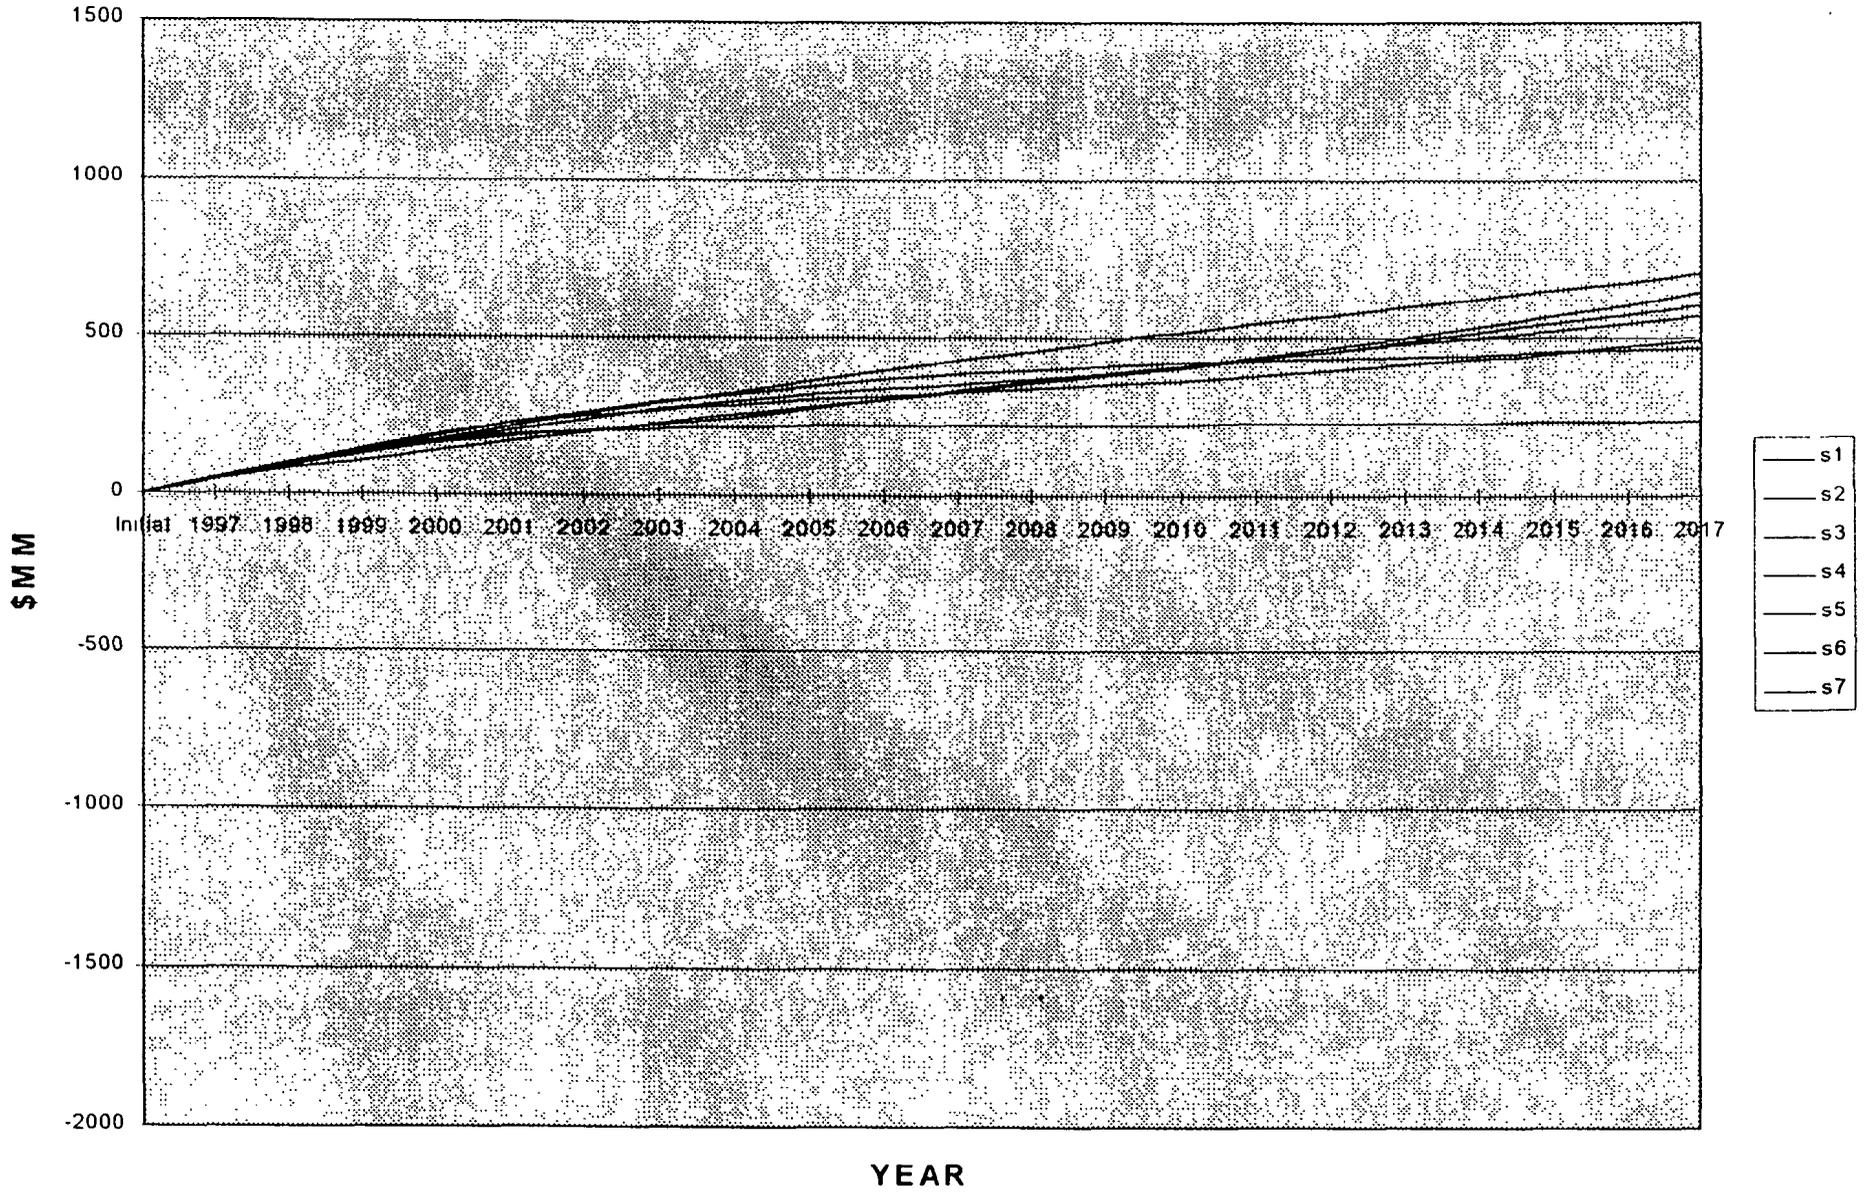
Chart 4 is proposed improvement and distributable earnings. This first bar demonstrates that in 1998, for this particular business line, we had a substantial improvement in distributable earnings largely due to the fact that in scenario seven, the down-rate scenario, there's a great deal of accrual that occurred. Many discount bonds got called and where this is an average debt, scenario seven was like the $45 million improvement. So this is a little bit skewed, but overall, we'll show you the individual scenarios. Chart 5 is the distribution. It is not present value. It is just gross dollars out into the future. In scenarios five, six and seven, we actually hurt the company's earnings a bit. This company had specified to us that it was acceptable for them to lose a little bit of surplus and a little DE in the down scenarios if we could show substantial improvement in the up scenarios. The up scenarios are almost five or six times the magnitude of the down scenarios. The company was happy with this. The important thing is it can be constrained to whatever a client wants.

Chart 6 is a smaller business line. Because seven was the negative scenario for this particular company, we were able to pick it up in Chart 7. We weren't able to make it stellar, but it's better than negative. When we spoke with the people at A.M. Best and we were flipping through some of this stuff, this was pretty relevant to them. They said, if you can keep people out of a ditch with this stuff, then we don't have to worry about ever coming in and necessarily downgrading because of mismatched assets and liabilities.

Another business line is actually one of my favorites. There was really no trouble in Chart 8 except you had a great deal of divergence out at terminal surplus. We did have a pretty big negative.

In Chart 9 you'll see that we picked up the negative, and more importantly, we shifted all the terminal surpluses. In every year of development, surplus was improved for this particular business line. Now we don't expect you to read the graphs discretely. What we generally do is prepare tables that show you all the data points and the differences and so forth.

For the company in Chart 10, the distributable earnings improvement was straight across the board with the exception of S-5 beginning in the year 2011. However, the company was pretty happy with that. Chart 11 shows more of an average type basis, but again, it showed overall improvement.

CHART 4
Proposed Distributable Earnings Improvement*
(TAS Output Treasury Reinvestment Assumption)

Present Value of Average Improvement weighted
equally across the NY 7 scenarios evaluated

PV Discount Rate

▨ 7.00%
▦ 6.00%
☐ 5.00%

$MM

YEAR

Initial 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017
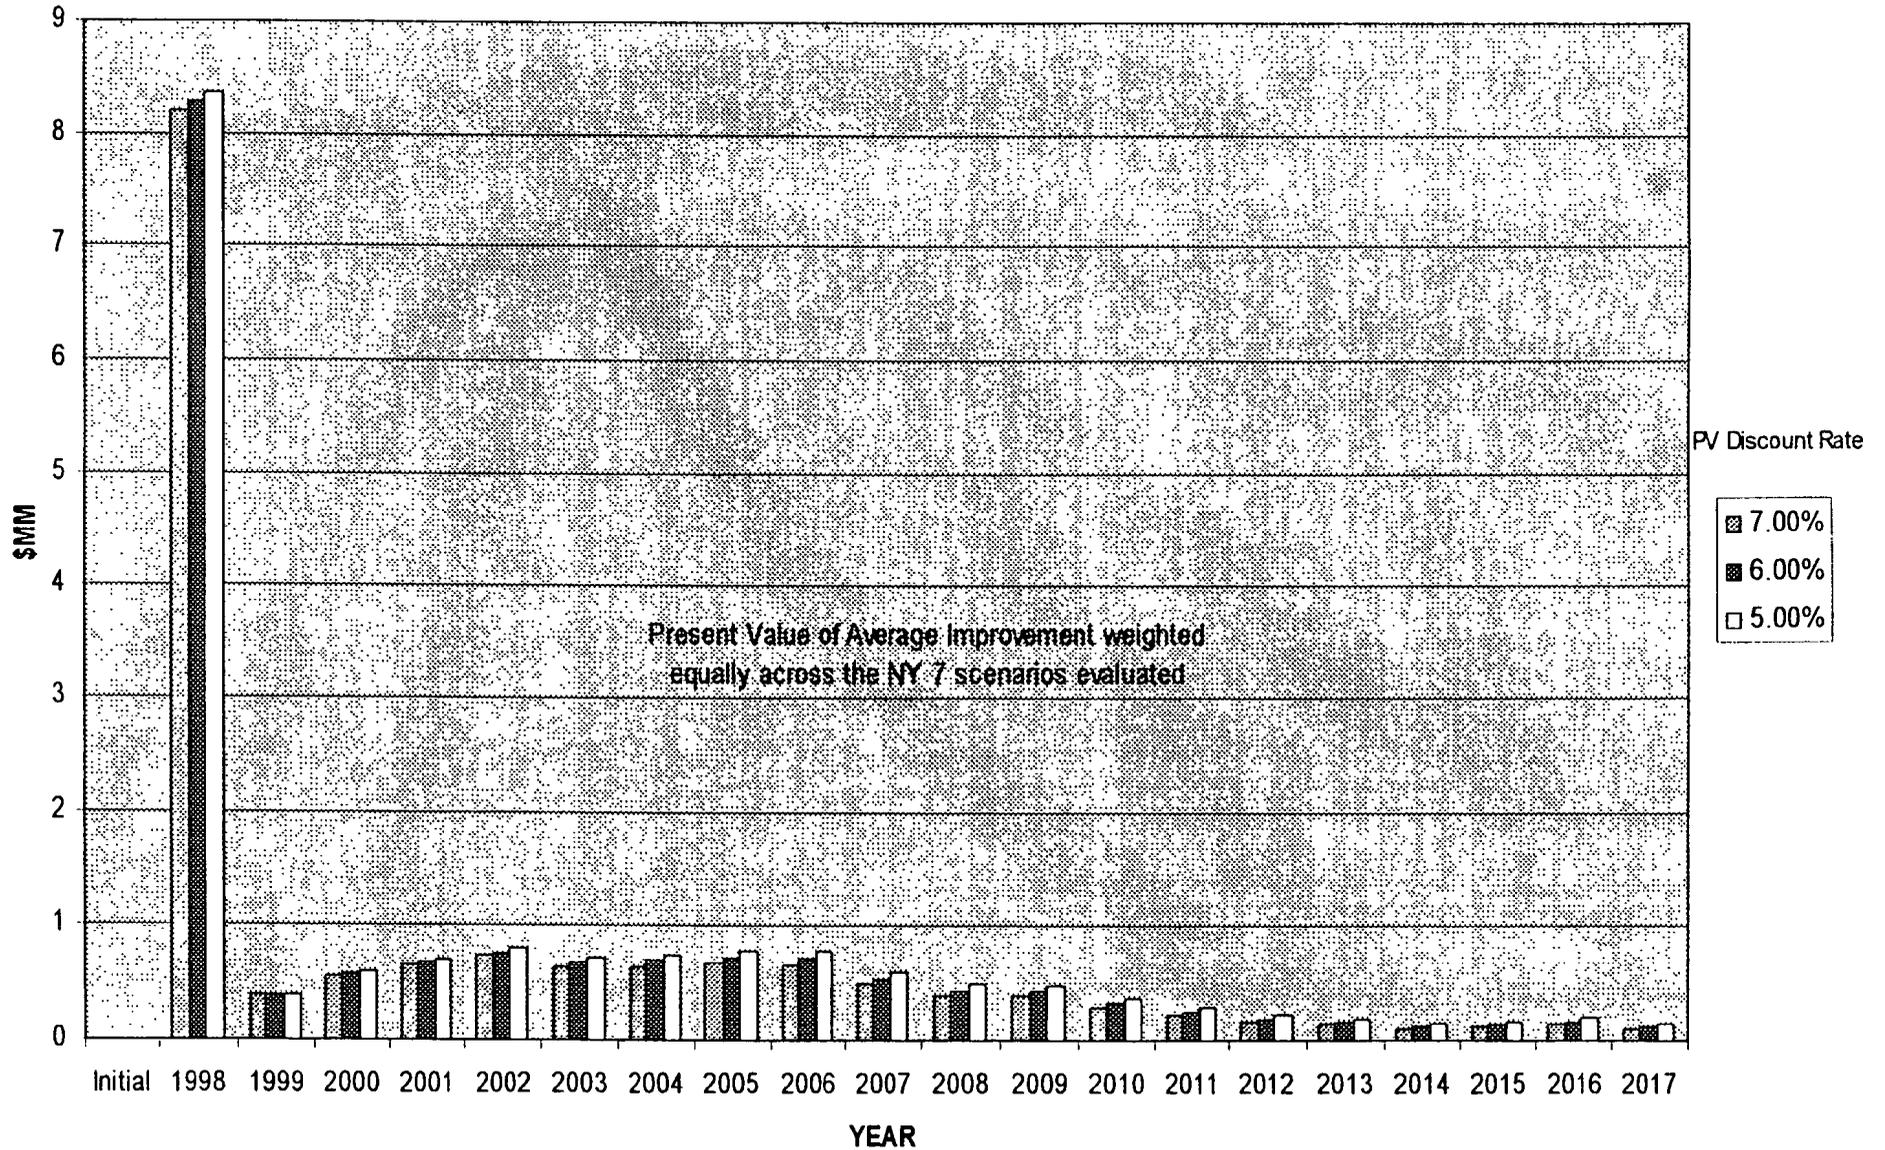
CHART 5
Proposed DE Improvement
(TAS Output Treasury Reinvestment Assumption)

Legend:
- S1
- S2
- S3
- S4
- S5
- S6
- S7

Y-axis: $MM (from -100 to 100)

X-axis: YEAR (Initial, 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017)
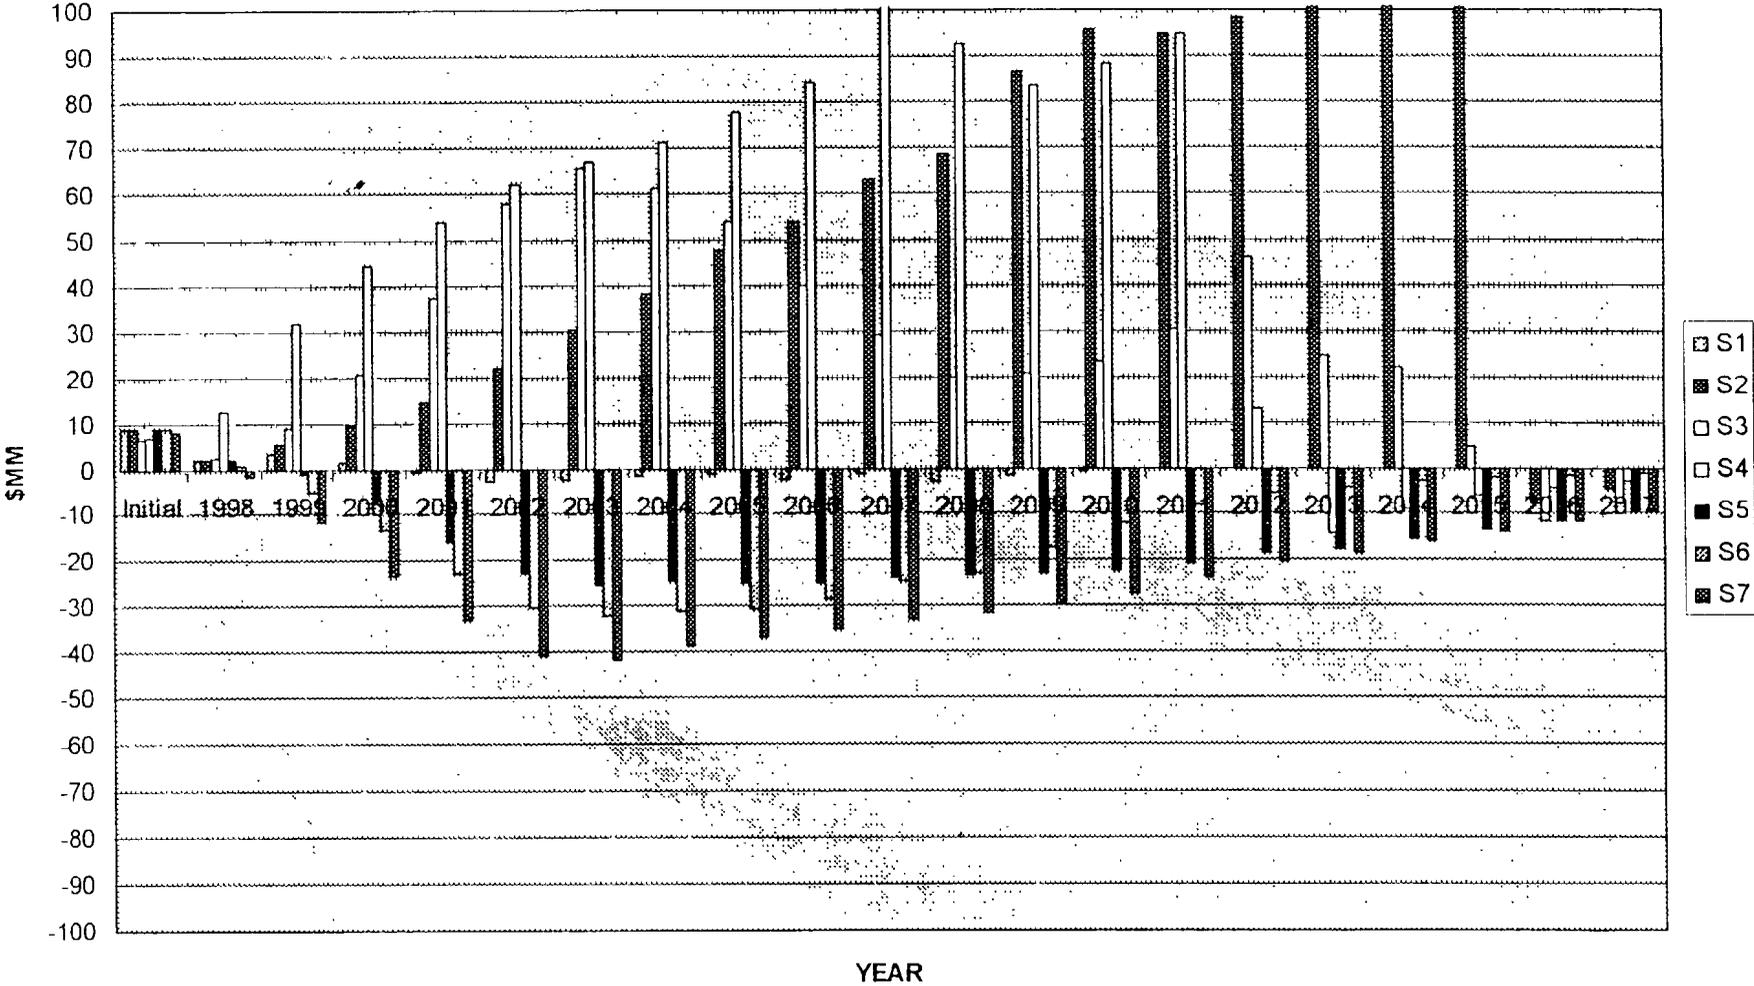
CHART 6
Current Surplus Development
(PTS Output Treasury Reinvestment)

YEAR

$MM
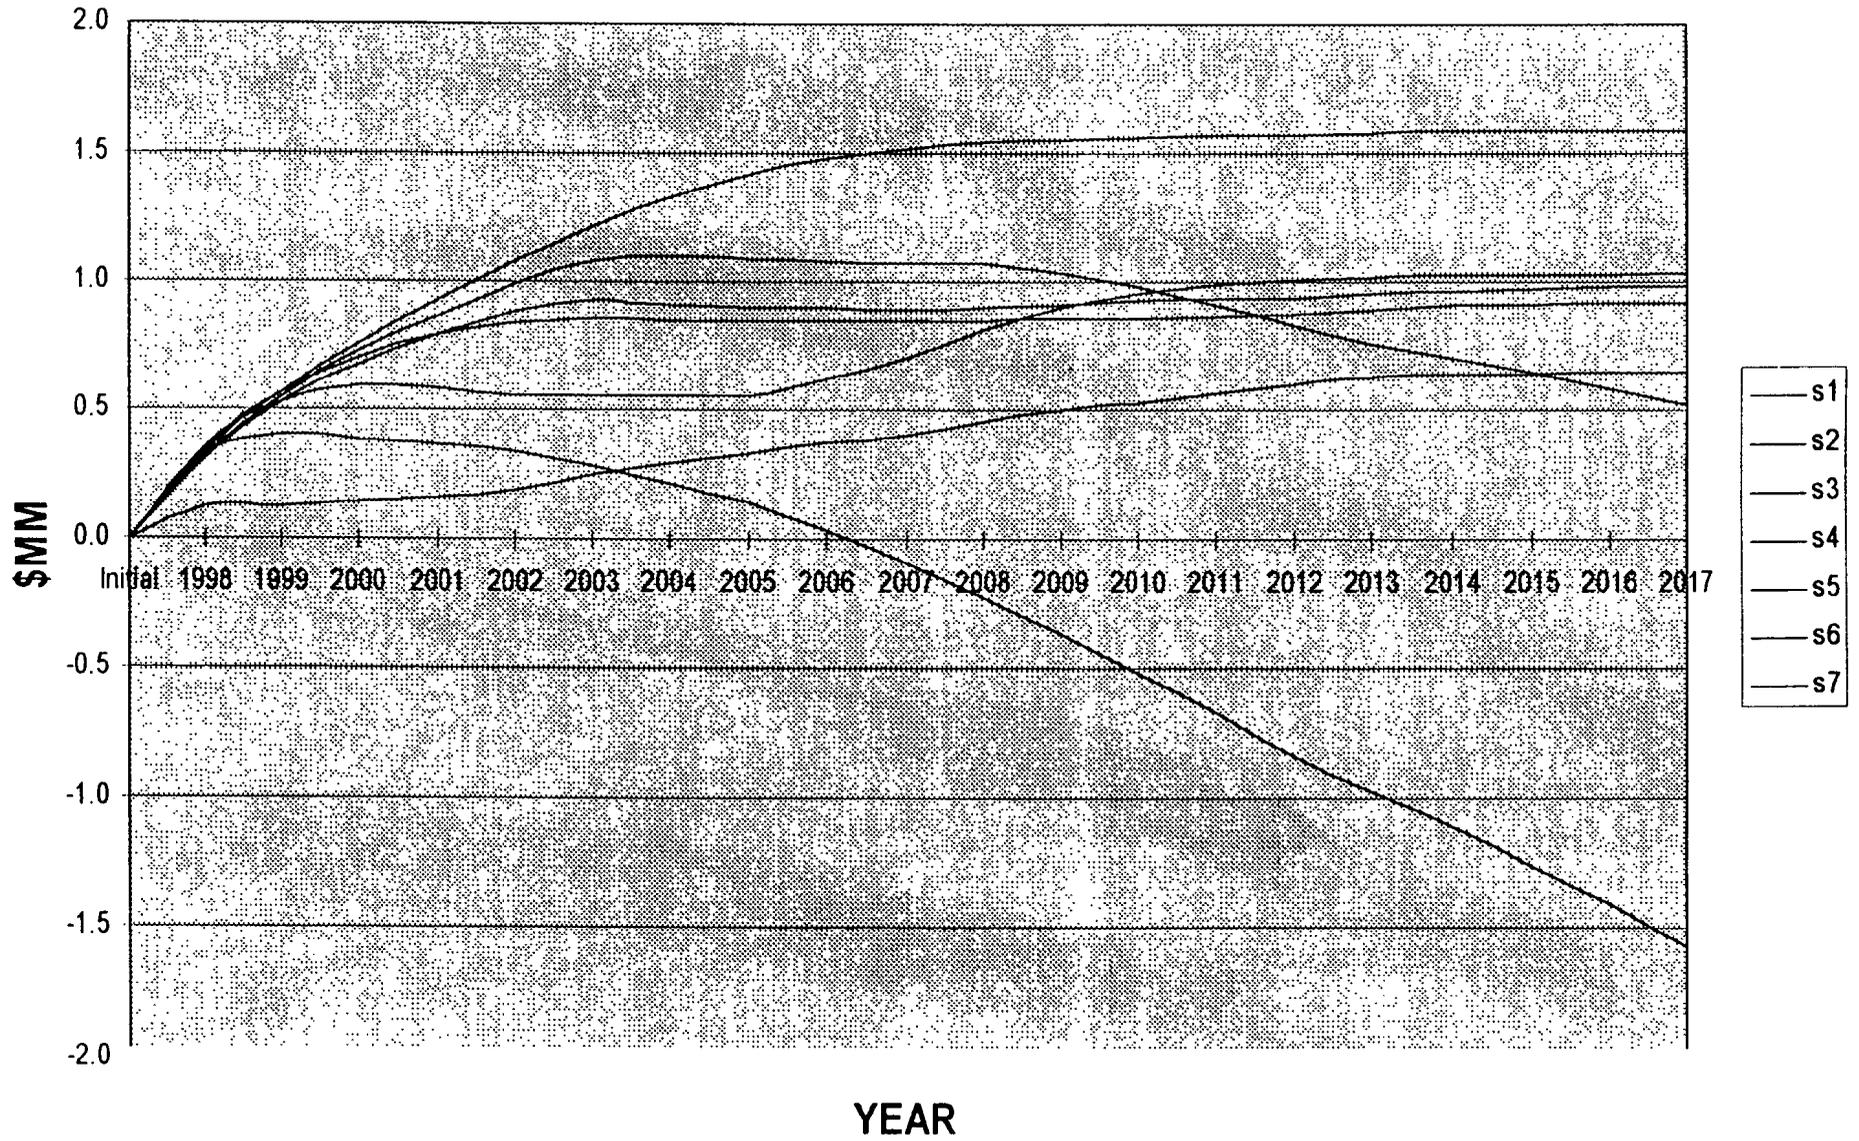
CHART 7
Proposed Surplus Development
(PTS Output Treasury Reinvestment)

Legend: s1, s2, s3, s4, s5, s6, s7

Y-axis: $MM

X-axis: YEAR (Initial, 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017)
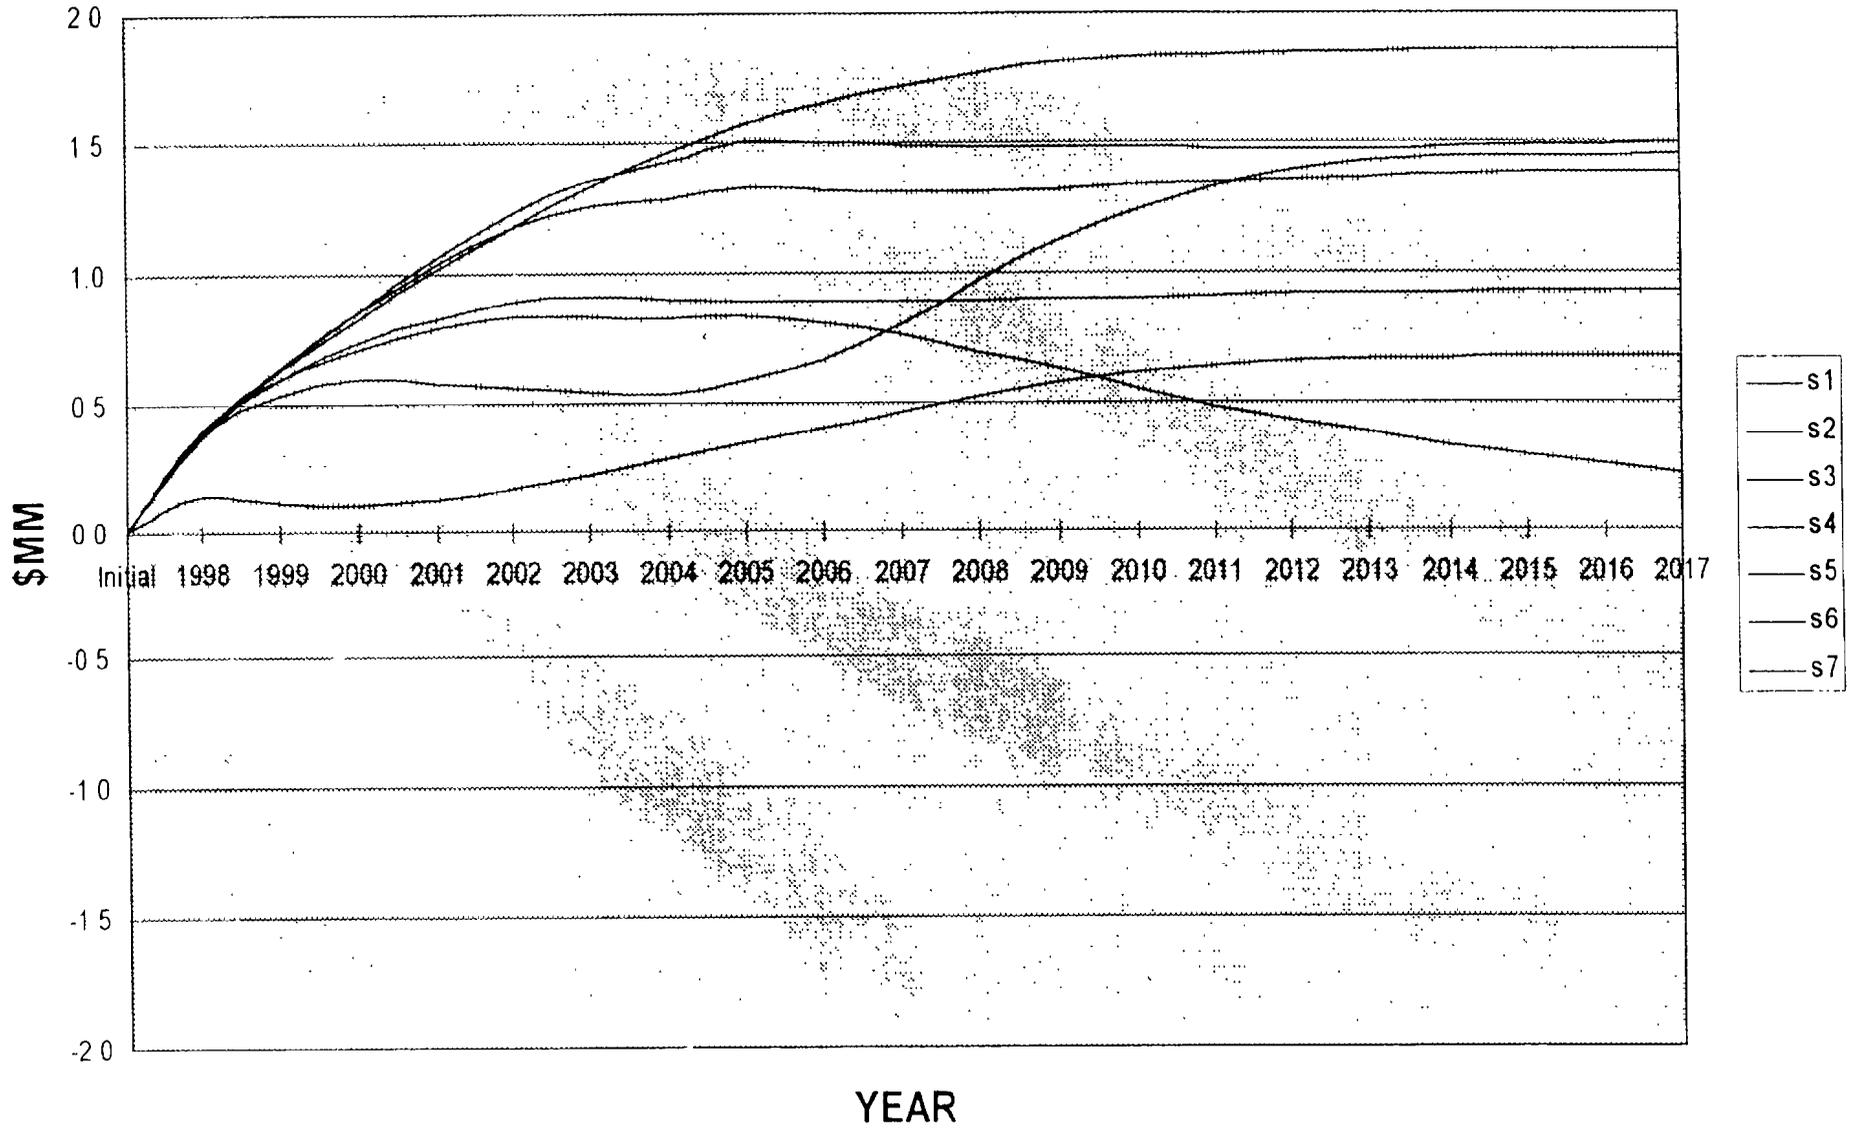
CHART 8
Current Surplus Development
(TAS Output Client Reinvestment Assumption)
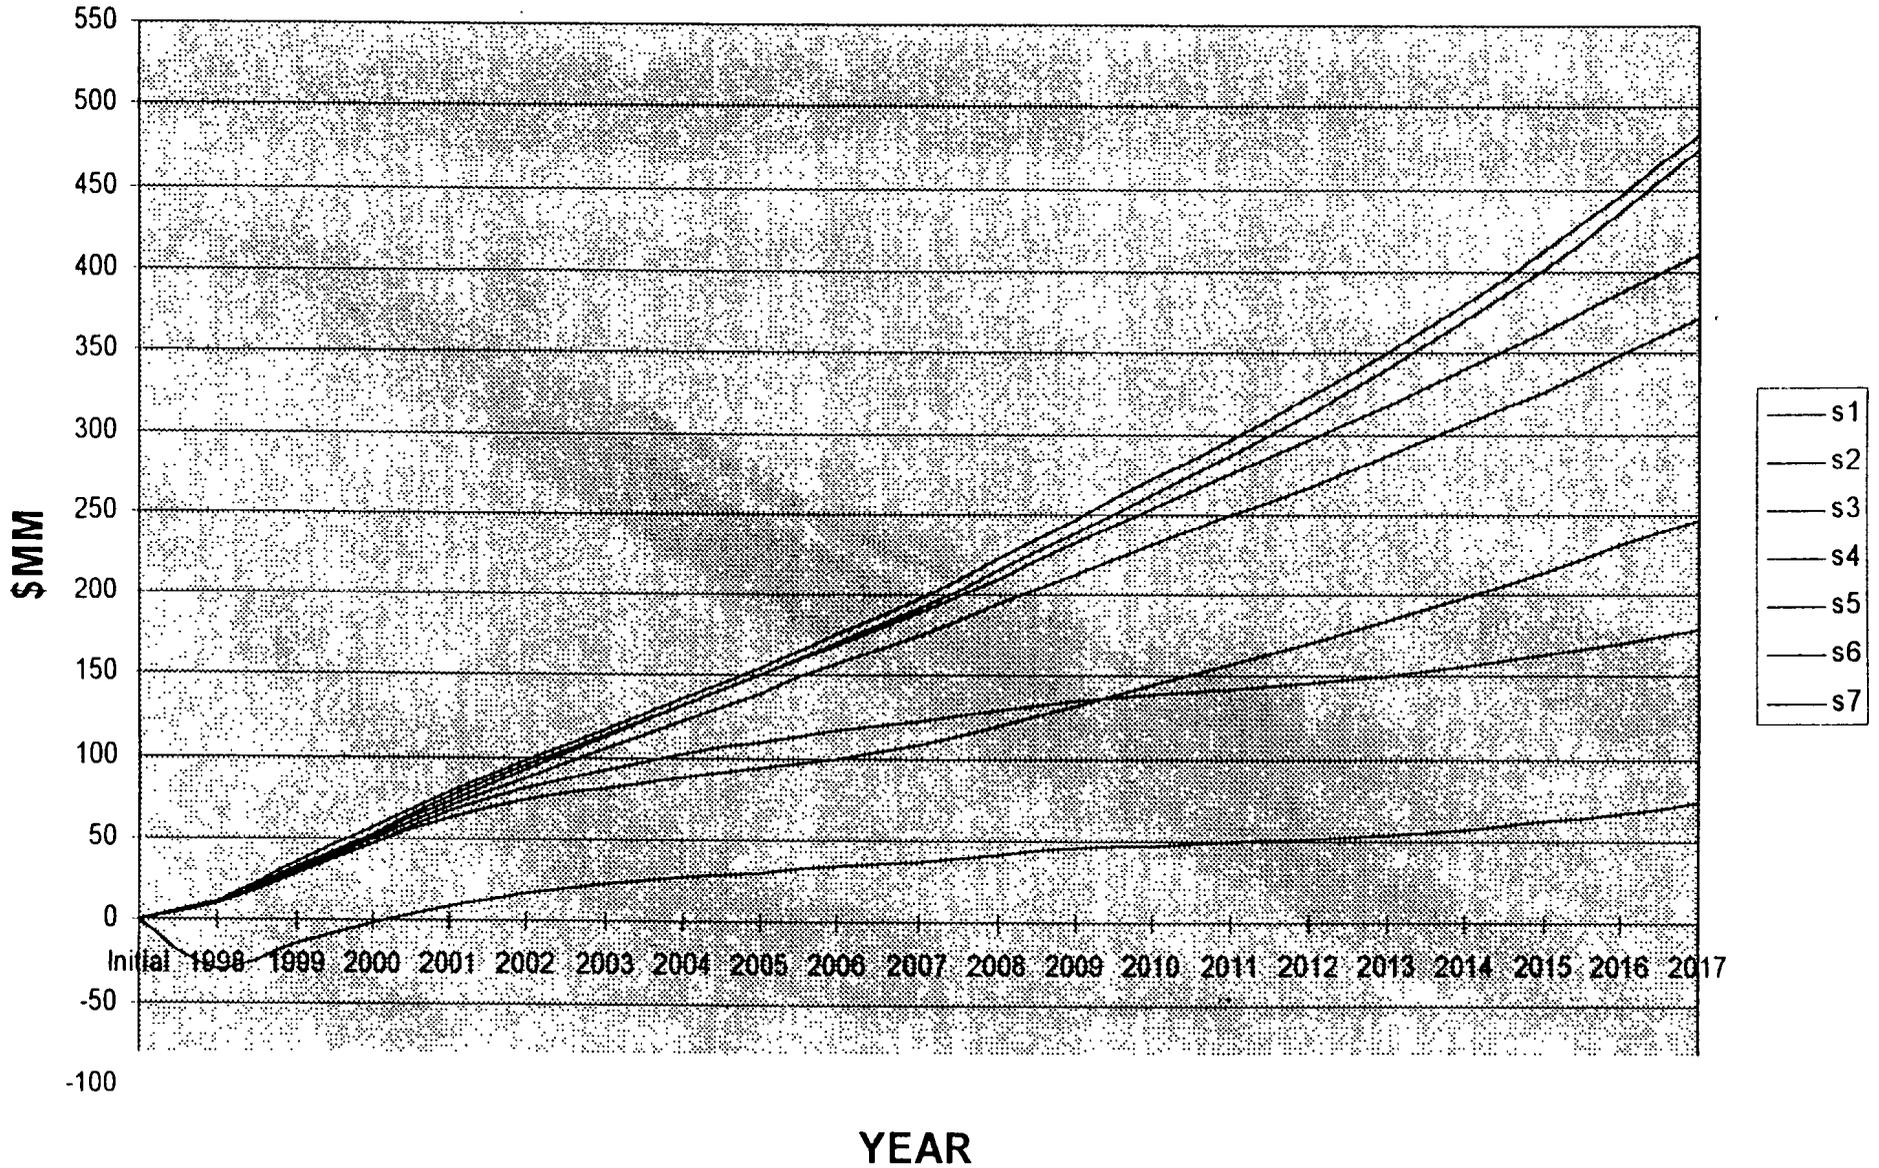
CHART 9
Proposed Surplus Development
(TAS Output Client Reinvestment Assumption)
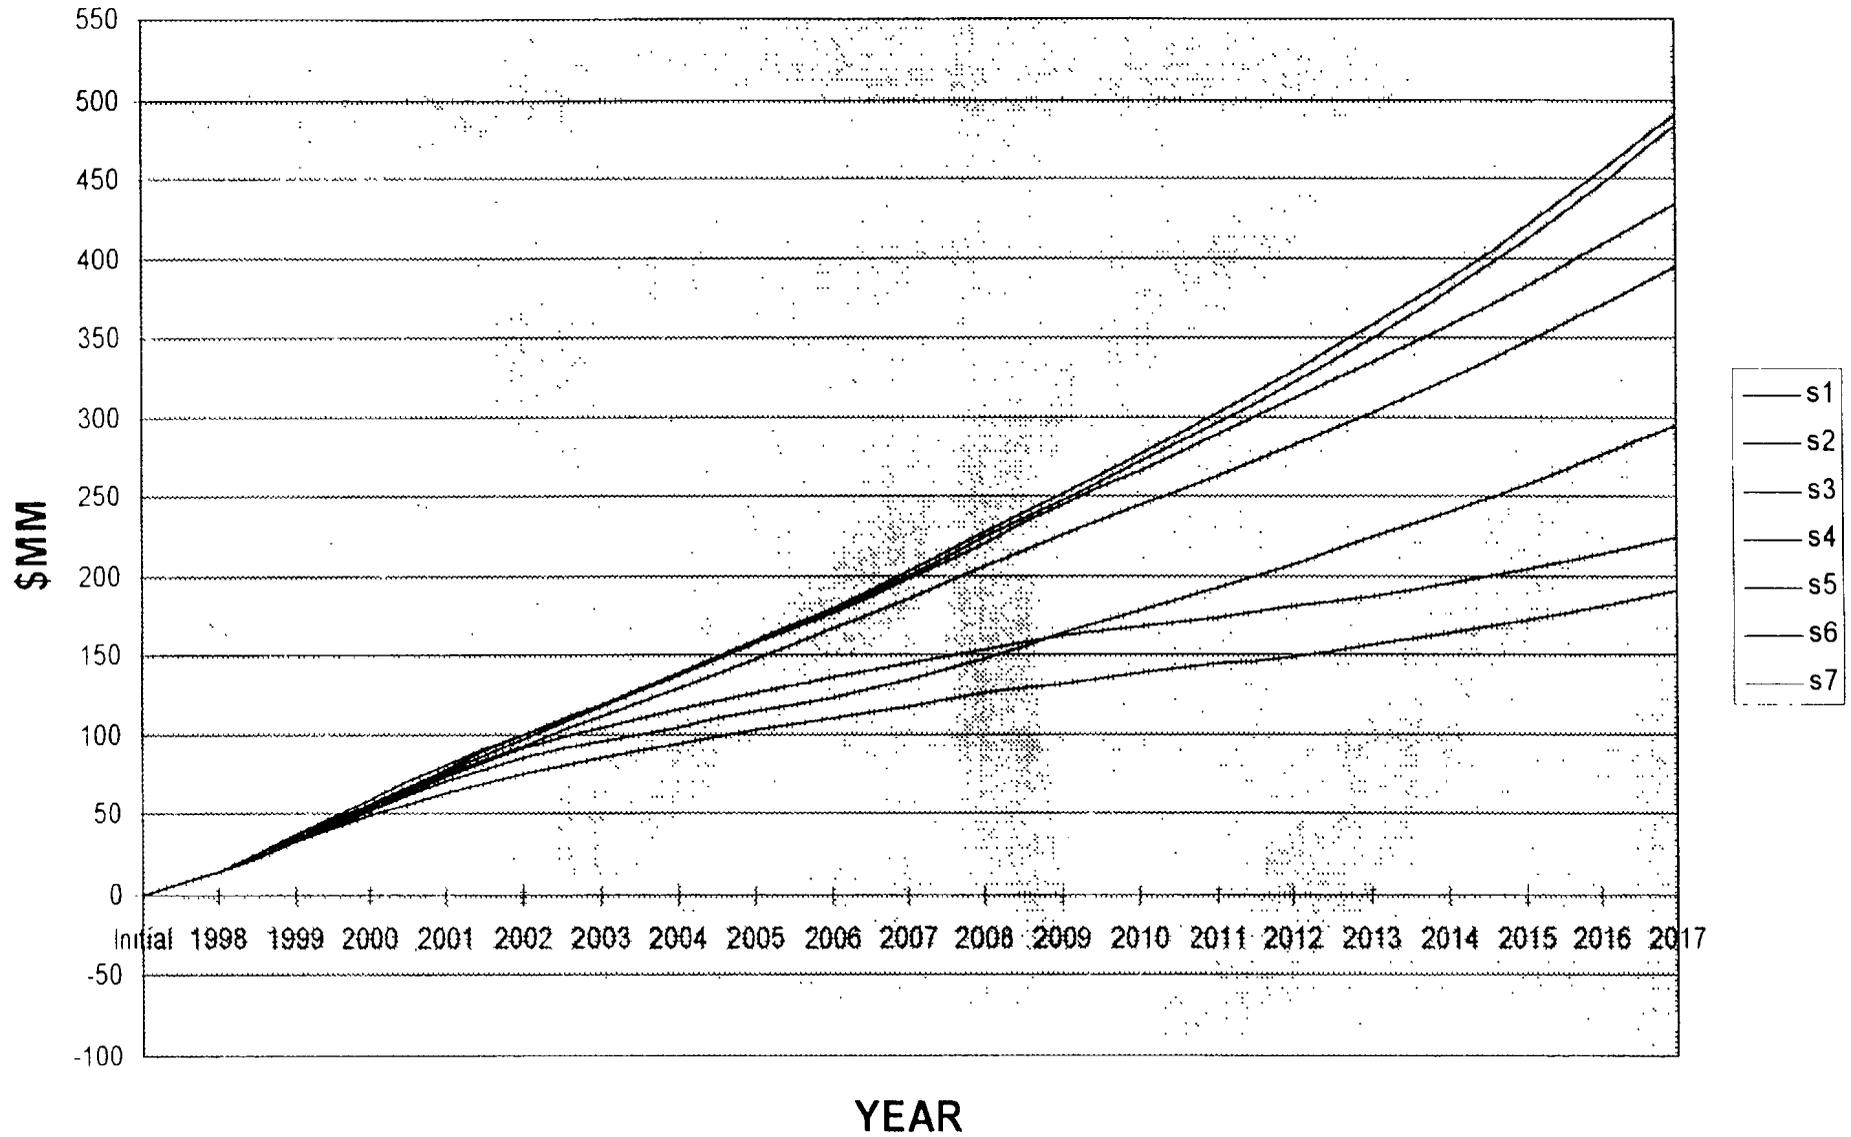
CHART 10
Proposed DE Improvement
(TAS Output Client Reinvestment Assumption)

Legend: S1, S2, S3, S4, S5, S6, S7

Y-axis: $MM (from -1 to 5)

X-axis: YEAR (Initial, 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017)
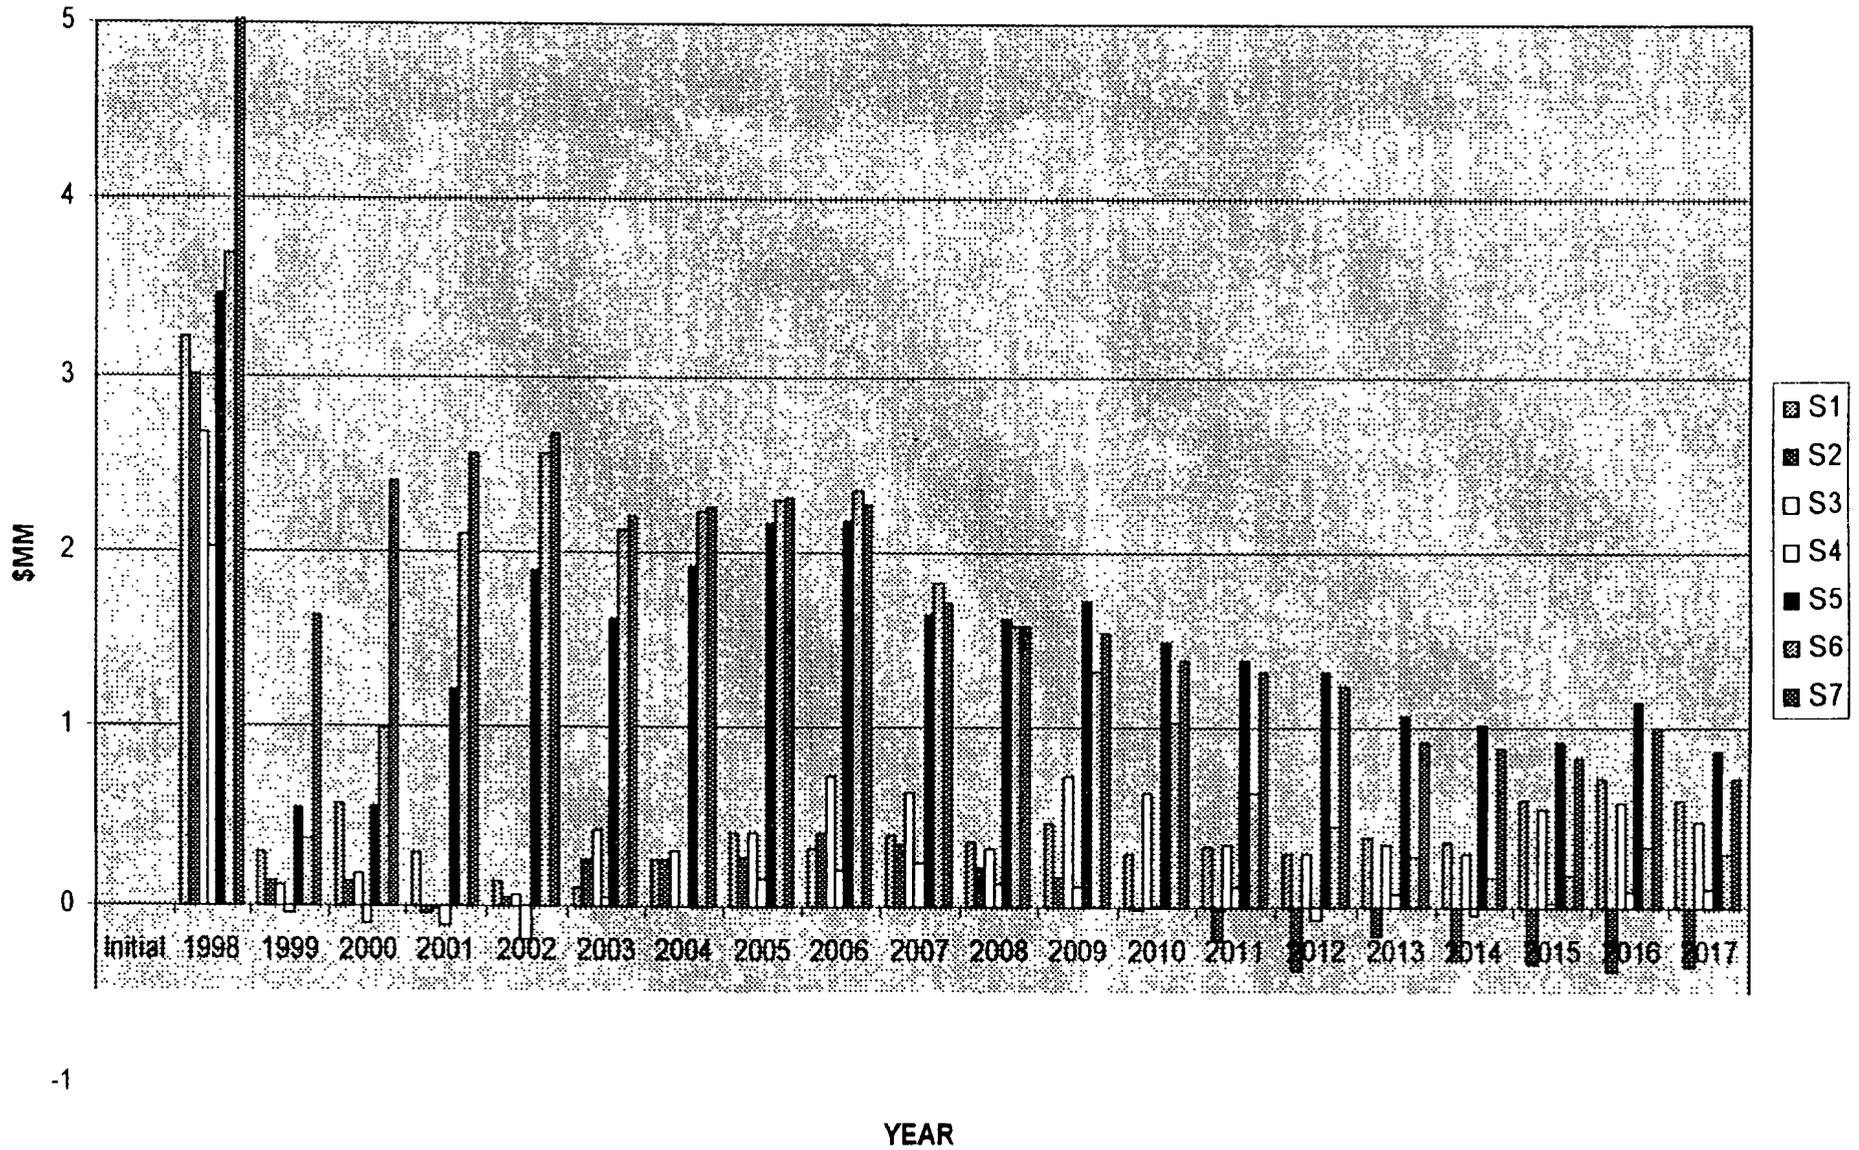
CHART 11
Proposed Distributable Earnings Improvement*
(TAS Output Client Reinvestment Assumption)

Present Value of A?erage Improvement weighted
equally across the NY 7 scenarios evaluated

PV Discount Rate

- ☑ 7 00%
- ◼ 6 00%
- ☐ 5.00%

$MM (y-axis)

YEAR (x-axis): Initial 1998 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017
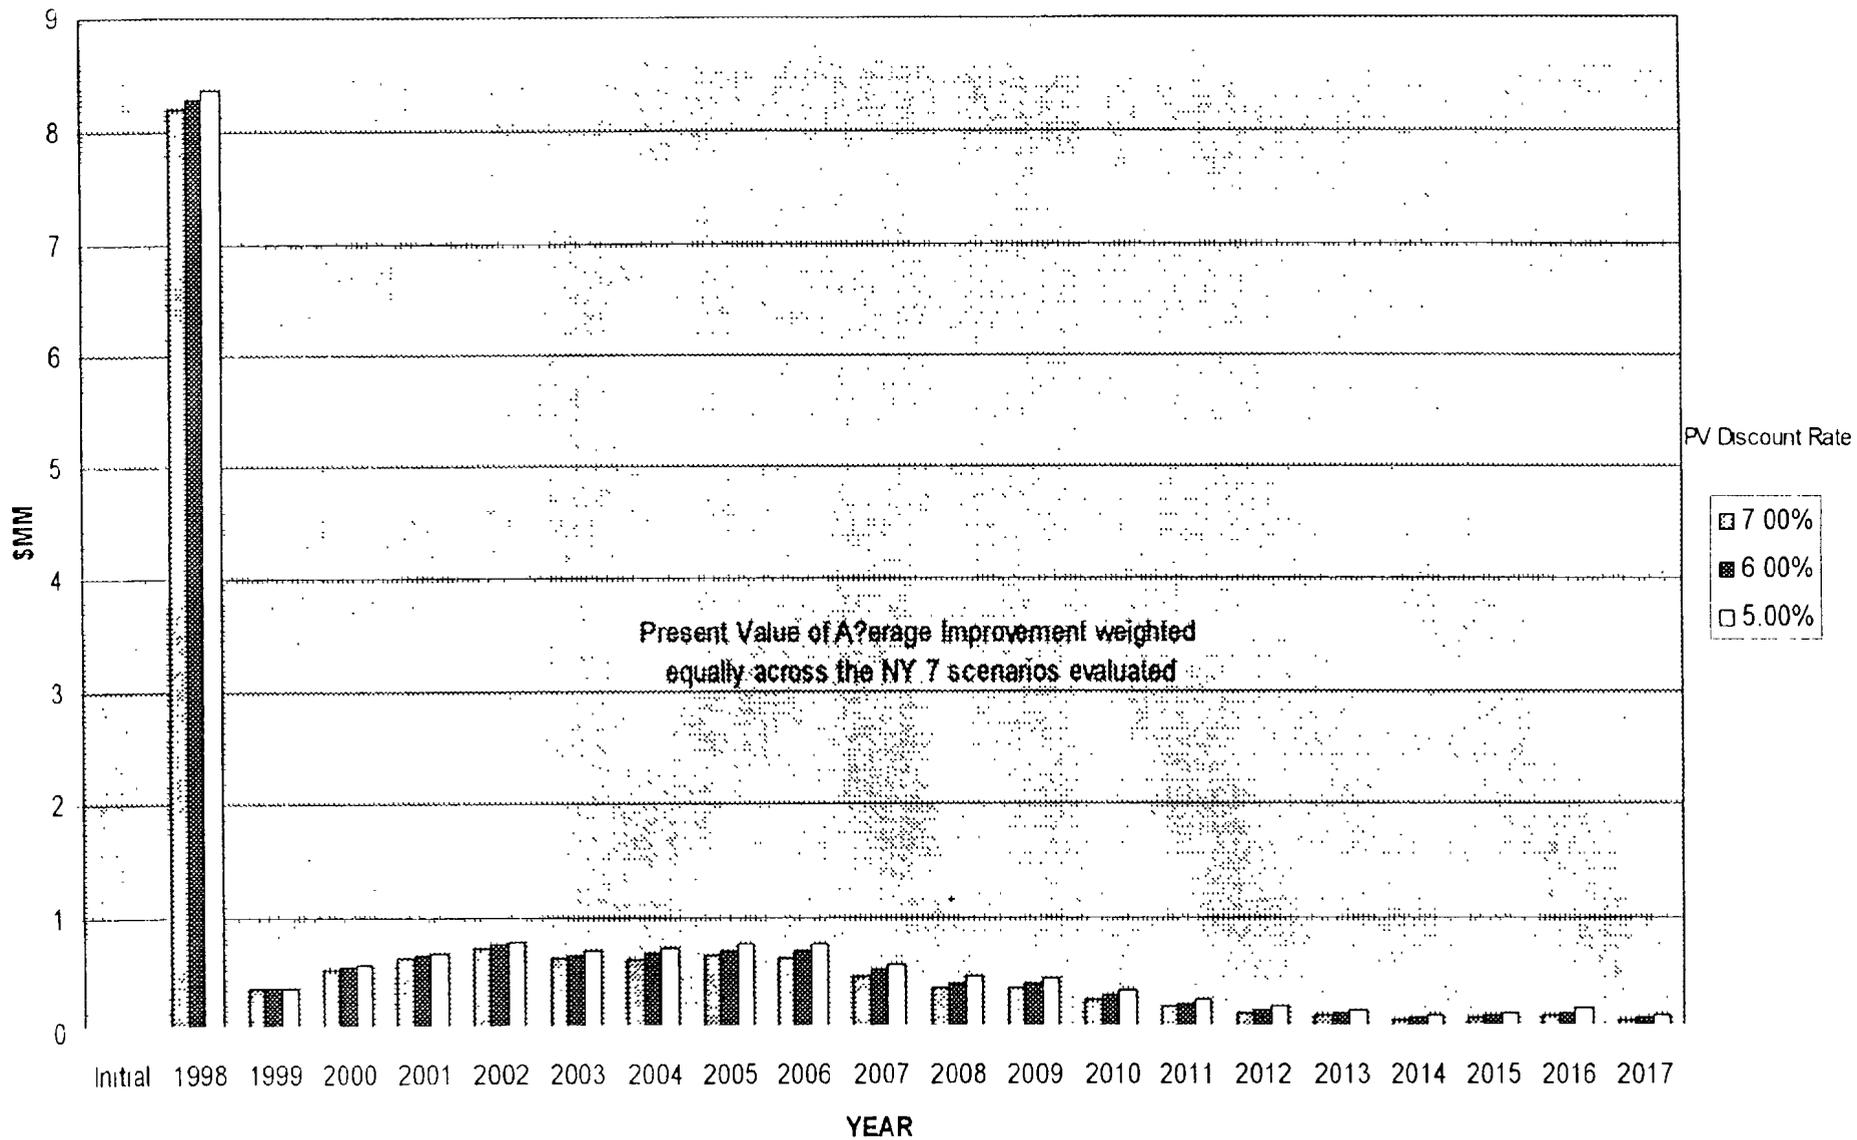
We are constantly asked what some of our constraints are. In fact, my partners are always saying we really need to write down all the constraints so we can tell people what we can do. The problem that we've had with that from day one is that there are really more constraints than we can write down, and we really don't want to suggest to anyone that there's a limitation because the program is fairly robust. It can accommodate a number of contingencies.

There has been a great deal of talk about efficient frontier immunization, which is super set of duration matching. You don't hear that said much, but that's the case. You hear cash-flow matching and you hear total return optimization. I'm familiar with each of these processes, and each of these processes has a limitation to it. Some work was done by some actuaries, *Elton and Gruber*, on efficient frontier analysis. They went through a very prolonged logic set that pointed out that for a company with specific liabilities, the efficient frontier boils down to a single point. I read their paper. If you think about it you'd realize you have specific liabilities. Does this total return paradigm really make sense for you because you have a liability in ten years? It doesn't matter how well you do today if you can't find it in ten years.

So what we have developed is cash-flow matching and total return optimization as constraints. Duration matching really didn't enter into it, but a cash-flow matched portfolio is a duration-matched portfolio, and all of these have been embedded in what we're doing as constraints. We specify as the objective function to make more money in some form or another surplus, and DE. Our investment *assets* are specific. We can allow negative assets. It allows you to form derivatives, if that would happen to fit, and it will accommodate proxy assets. It's dumbing down the system a bit. We can accommodate any derivative for which you can give us cash-flow, book, market and accrued income values.

Other constraints are: beginning cash, diversification, sector as well as overall total transaction size, taxation, and scenario specification. Borrowing is a limitation in our system due to the fact that we deal with a *solver*. You allow the sale of assets as opposed to borrowing because if we allowed assets sales, the problem would become nonlinear. This is undesirable. Moreover, a solver will exploit arbitrage opportunities. Most companies cannot deal with such risk.

Yield is calculated as required. As for reinvestment, our routines are set up to be virtually identical to *TAS* and PTS and *ALFA* with the one exception. We do not feel comfortable with modeling forward collateralized mortgage obligation (CMO) cash-flows as reinvestment type issues because we think that it would require a complete wild guess. We think that it may look good on a cash-flow test or in this type work, but we really don't think it's valid. Other constraints are *hold to maturity vs. available for sale*, surplus targeting and other constraints that can be added as mathematically specified by the client.

The securities universe is really anything our client wants it to be. Now we subscribe to a service where we see about 5,000 securities a day. Of those, as you might be able to imagine, a number of them are duplicate securities, some of them are illiquid securities, and some of them are Korean or Indonesian type securities that we have been chastised for lately so we've taken those out. Once we surveyed Wall Street firms, we normally have 2,000–3,000 securities plus what's already in the portfolio, plus anything else the client wants to come up with. In many cases, clients have said, "Send us your universe." There are certain securities that we'd never buy because it's against our policy, and we'll take those out. We have no axes to grind. We're not brokers, and we're not trying to sell securities. We don't care where your universe comes from. Consequently, Treasury notes STRIPS, corporates, agency notes, mortgage collateral, CMOs, and client-specified assets are all fine. Whatever you want to put in there will work.

Let's discuss a typical portfolio restructure (Table 1). I love this table because someone will always walk up to me after I finish speaking and say, "Our portfolio is already sectored like that. I bet you can't improve it for us." I spoke at the ALM conference in New York, and a follow came up to me and said that I said, "Why don't you give it a try," and he did. The company had a duration of five-and-a-half years. It actually had this sort of structure, but its specific liabilities called for something else. We showed substantial improvement in every scenario as compared with this particular company. We had almost exactly the same duration that they did. They wanted to know what the deal was with the duration. If you haven't changed it, how can you get better results? The thing is, and this will appeal to the mathematicians in the group, duration is not a unique calculation. There are many solutions that will fit for a given duration. What's more important in immunization is the

cash-flow disbursement portion of it is very significant, and that's what we're capturing. On top of that, we're specifying that we want to make more money. In one of the case studies that I showed, the company was 65% into corporates to begin with. They went to 74%. They went from 20% down to 11% for marketable mortgages. Private placements were constrained because we didn't really feel like we had a good handle on the liquidity. Duration changed a bit. We typically find that the duration goes out a bit for no specific reason. Convexity is almost a nonissue.

TABLE 1
Typical Portfolio Restructure

	Before	**After**
Corporate, Treasuries, Agencies	65%	74%
Marketable Mortgages	20%	11%
Private Placements (constrained)	15%	15%
Duration	3.64 yr	4.50 yr
Convexity	−0.05	+0.06

Table 2 shows the output. This is the optimal portfolio. This is actually what we call a targeting grid. To read it, you must begin over at the far right side. You have a specific *CUSIP* (we made these up). This process assumes that at time zero, this new portfolio is in place. That never happens all at once or immediately. However, this can be used to target as you go forward because asset managers can target in this fashion. When we show another report, which is literally the optimal portfolio, we have the positions of the securities.

Look at the top security for example, the price, the coupon, the maturity, the spread, the duration, the convexity. If the portfolio manager does not want to buy that particular *CUSIP*, he or she is certainly capable of going out to the street and buying something that's a lot like it. If they really want to dumb the process down, they can do that and they can go to the duration and convexity numbers and things that they're comfortable with. However, this gives them so much more

TABLE 2
Optimal Solution

CUSIP	DESCRIPTION	CURRENT POSITION	PROPOSED POSITION	DIFFERENCE	PRICE	COUPON	MATURITY	SPREAD	DURATION	CONVEXITY
123456AB25	Specific Security Name	10,000,000	15,000,000	5,000,000	100.40	6.77	1/19/06	0	4.92	-0.74
123456AB26	Specific Security Name	10,000,000	15,000,000	5,000,000	101.43	7.56	6/12/06	40	5.30	-0.26
123456AB27	Specific Security Name	10,000,000	5,833,333	(4,166,667)	105.60	7.83	7/26/06	45	4.08	-0.59
123456AB28	Specific Security Name	10,000,000	9,121,525	(878,475)	99.95	7.29	8/28/06	30	5.48	-0.06
123456AB29	Specific Security Name	5,000,000	5,000,000	0	102.00	9.38	1/1/98	35	0.88	0.01
123456AB30	Specific Security Name	10,000,000	0	(10,000,000)	101.50	7.25	5/15/99	55	2.06	0.05
123456AB31	Specific Security Name	10,000,000	0	(10,000,000)	96.11	5.63	2/1/01	40	3.42	0.14
123456AB32	Specific Security Name	10,000,000	0	(10,000,000)	102.28	7.25	8/1/02	0	4.34	0.24
123456AB33	Specific Security Name	5,000,000	15,000,000	10,000,000	102.00	7.75	1/15/99	30	1.79	0.04
123456AB34	Specific Security Name	10,000,000	0	(10,000,000)	99.79	7.10	3/15/06	65	6.47	0.54
123456AB35	Specific Security Name	5,000,000	0	(5,000,000)	110.18	9.57	4/15/01	45	3.42	0.15
123456AB36	Specific Security Name	10,000,000	15,000,000	5,000,000	101.50	7.20	11/1/06	45	6.84	0.6
123456AB37	Specific Security Name	5,668,000	0	(5,668,000)	110.25	9.25	10/1/05	55	4.00	0.1
123456AB38	Specific Security Name	0	0	0	103.39	9.38	3/15/98	0	1.03	0.02
123456AB39	Specific Security Name	0	0	0	105.75	9.13	8/1/99	20	2.14	0.06
123456AB40	Specific Security Name	0	15,000,000	15,000,000	107.60	9.63	7/15/00	25	2.94	0.11
123456AB41	Specific Security Name	0	15,000,000	15,000,000	100.30	6.80	5/22/01	25	3.19	-0.42
123456AB42	Specific Security Name	0	15,000,000	15,000,000	102.30	7.50	1/15/02	15	4.10	0.2
123456AB43	Specific Security Name	0	15,000,000	15,000,000	103.60	7.63	1/15/07	55	6.95	0.62
123456AB44	Specific Security Name	0	0	0	100.00	7.88	3/31/02	60	3.50	-0.54

information that they can use to do their asset targeting. At any rate, this gives much more specific information that can be passed on to an asset manager than things like covariances and durations and convexities. Because I have traded, I know I'd much rather have a road map like this than a few benchmarks.

Chart 12 is our workflow diagram, which shows how we go through the process. We periodically get asked the question: "If you change the asset portfolio, don't you, by definition, change the crediting rate? Don't you, by definition, run off more of the liabilities a little sooner, and doesn't that unbalance the relationship?" The answer is, yes, it does. We have developed a few techniques to freeze the liabilities a little better, but ultimately they do change. If we have to do one or two iterations of this process, the liabilities and the assets will converge and that's how we deal with particularly volatile liabilities. However, we have found that most of the liabilities that we were working with weren't nearly as volatile as the embedded optionality of most of the assets. So we still feel like the asset side is the wild card here. It's the one you can do the most about largely because you're in a market. You're trying to sell into that market, and the market dictates much of what goes on on the liability side.

CHART 12
The AQS Process–Workflow Description
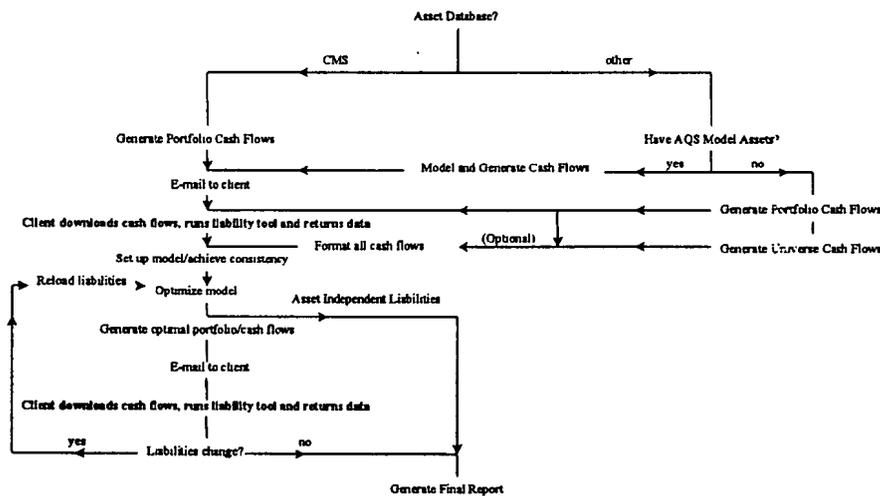
MR. MEL STEIN: I'm going to go over an overview of the stochastic analysis area. Stochastic projections are the most rapidly evolving areas of actuarial practice. I will cover the various aspects of stochastic projections including a number of new and exciting applications. Let's start with the planning of various aspects of stochastic runs like tracks of runs, types of scenarios, and types of analysis. Then I'll introduce some exciting new concepts such as the meltdown point and discuss critical state-of-the-art applications.

Let's discuss types of stochastic runs. First are the single-point runs that most of you are used to. This is a single strategy or set of strategies for a single product line or strategic business unit. Next are the vertical stochastic projections where you project the entire company or group of companies. It sort of looks like a pyramid. The purpose is to evaluate risk for the entire operation and project the performance of the entire operation. Last, are the horizontal stochastic projections in which you simultaneously test a number of alternative strategies at the same time for a specific strategic business unit in a single run. For example, you can test six or eight or ten crediting strategies or investment strategies.

In terms of types of stochastic scenarios, there is interest only, which is what all of you are used to. Next is the first type of dual dimension run, the interest rate-stock performance scenario. Ideally it would have at least 10,000 scenarios (a hundred interest times a hundred stock performance scenarios). This is applicable to variable and equity-indexed products. Companies with these products should be doing this, and some of them are. Last, we have something brand new: the interest rate state-of-the-economy projection. Again, this would ideally have at least 10,000 scenarios (one hundred interest times one hundred state-of-the-economy scenarios). All companies should be running these. Very few, if any, are.

Now let's look at the evolution of ALM scenarios. The interest rate state-of-the-economy stochastic projection is a critical step forward in this evolution. For many years, we used the single-scenario asset-share type of projection. Then we projected assets as well as liabilities, and this led to the amazing discovery that interest rates go up and down. The next step was the stochastic projection based on 100 or more randomly generated interest scenarios. This was necessary to define the

interest rate risk. The interest rate state-of-the- economy stochastic runs are necessary to recognize the economic, noninterest-rate risk. The July 1988 draft of the actual principles of ALM from the ALM Principles Task Force states that an effective ALM process consists of five fundamental steps, the first of which is identifying the level of exposure risk. We have been doing this for interest rate risk, and we have not been doing this for economic risk. We have been incredibly lucky in the 1990s.

Why are 10,000 scenarios required? Why do we need 100 economic scenarios for every interest scenario? Because the scenario variables include: when the bad economy starts; how bad is bad; whether it is a depression, a major recession, or a moderate recession; the state-of-the-economy that precedes it; and the state-of-the-economy that follows it? Does the bad economy occur with high, rising, dropping or low interest rates? How does this change during the bad economy? How is each state-of-the-economy defined? What is a major recession? The answer is defined by its effect on assets and liabilities, such as, defaults and policy loans. A bad economy is bad because of its effect on insurance company operations.

How is the state-of-the-economy projected? We use a Markov chain approach where the user controls the probability of the state-of-the-economy not changing or changing to any one of the other states of the economy. The user is allowed to specify the state-of-the-economy for the first X time periods, and then have the Markov chain take over generating subsequent scenarios. We do not use an economic model to forecast the economy. All this does is add a bunch of additional assumptions that are likely to be no better than SWAGs. In other words, you add an element of garbage in, garbage out (GIGO). Furthermore, if you put five economists in a room, you might very well get five different opinions or forecasts. Finally, economic models are, for the most part, for short-term forecasts. The question is not why the bad economy occurs, but, instead, when it occurs and how long will it last? This is what determines the effect the bad economy will have on insurance companies.

Let's go look at one more aspect of stochastic analysis—types of output. First, we have reports versus graphs. Graphs are normally invaluable in providing a clear, easily understood picture of critical data or comparisons. However, there are times when tables and numbers are equally valuable. Ideally, you should have both. Here are several ways of evaluating performance: (1) point-in-time-values, such as ending surplus or present value of earnings; (2) rate-of-return measures, such as ROE; (3) analysis over time graphs showing performance and trends over time; (4) probability analysis, such as probability of realizing financial goals; and (5) comparative risk/reward analysis where both risk and reward are defined both in terms of criteria and numerical values.

Stochastic output normally uses mean values, and quite often plus and minus one or more standard deviations, as well as the highest and lowest values. Chart 13 is a point-in-time graph showing ranked values of ending surplus. The number and level of the worst case scenarios is also shown.

Chart 14 is an analysis over time graph which shows the mean surplus, plus and minus one standard deviation, and the highest and lowest values of surplus at the end of each projection year. Again, we see the trends over time.

Chart 15 is a special type of point-in-time value graph. This shows the distribution of the values of economic surplus where the horizontal scale shows the value of economic surplus, and the vertical scale shows the number of scenarios that have a value of economic surplus within a certain range. This approach values both assets and liabilities by using spot rates to discount future cash flows.

Let's look at specific applications. Stochastic optimization of interest crediting strategies should result in a significant value in terms of mean expected financial results and/or reduction of the downside risk. Investment quality risk analysis becomes possible with real world risk analysis, which recognizes economic risk as well as interest rate risk. Optimal hedging strategies can only be accomplished through stochastic ALM runs. Effective duration and convexity of assets, particularly liabilities, are wanted by many of the larger investment departments. Both types of dual-dimension stochastic runs provide an effective risk analysis for variable and equity-indexed products. Now let's get into stochastic interest crediting.

CHART 13
Test

11.5M

9.5M

7.5M

5.5M

3.5M

1.5M

-0.5M

-2.5M

-4.5M

-6.5M

-8.5M

51.0
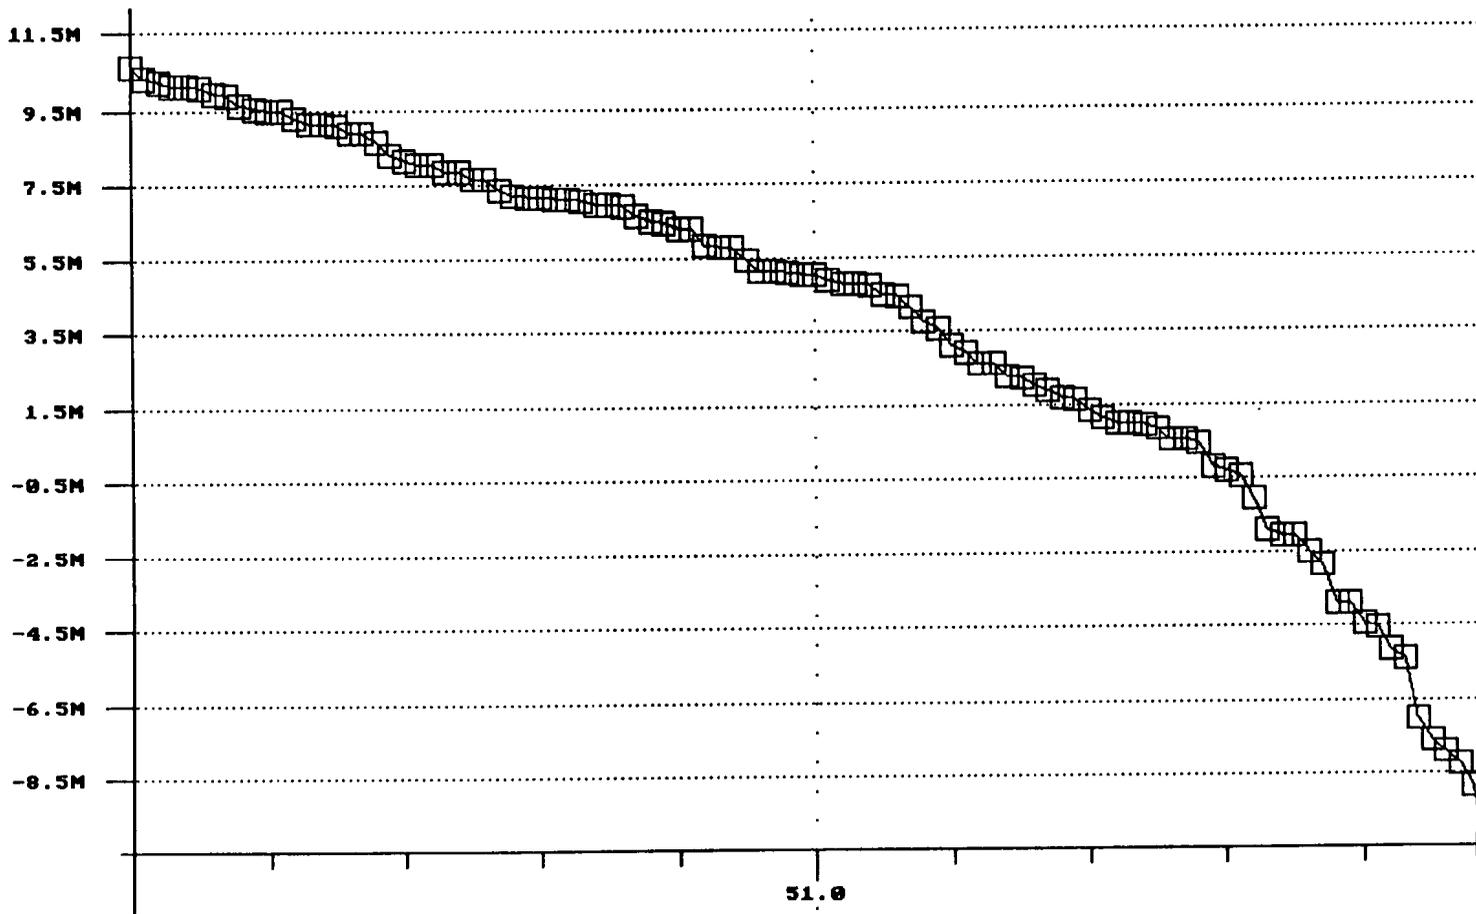
Ranked Scenario

☐ Ending Surplus

CHART 14
Test—Surplus

11.5M
9.5M
7.5M
5.5M
3.5M
1.5M
-0.5M
-2.5M
-4.5M
-6.5M
-8.5M
-10.5M

2.0 3.0 4.0 5.0 6.0 7.0 8.0 9.0 10.0

Year
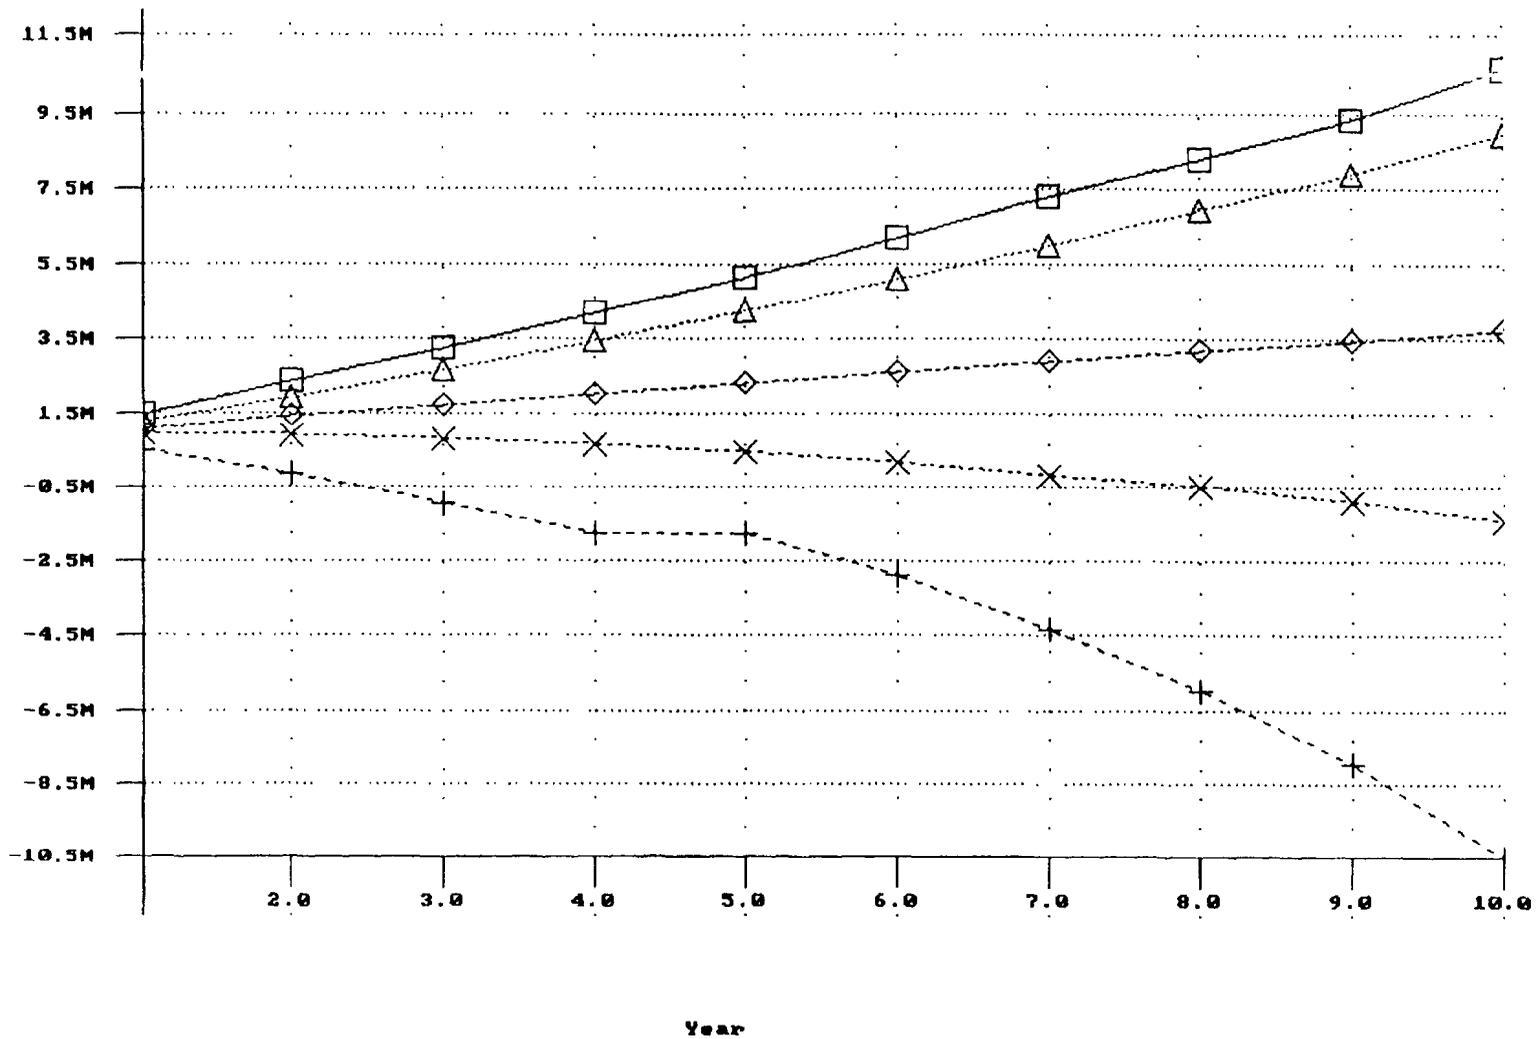
□ Maximum   △ m+StdDev   ◇ Mean   ✕ m-StdDev   + Minimum

CHART 15
Economic Surplus: PV Surplus
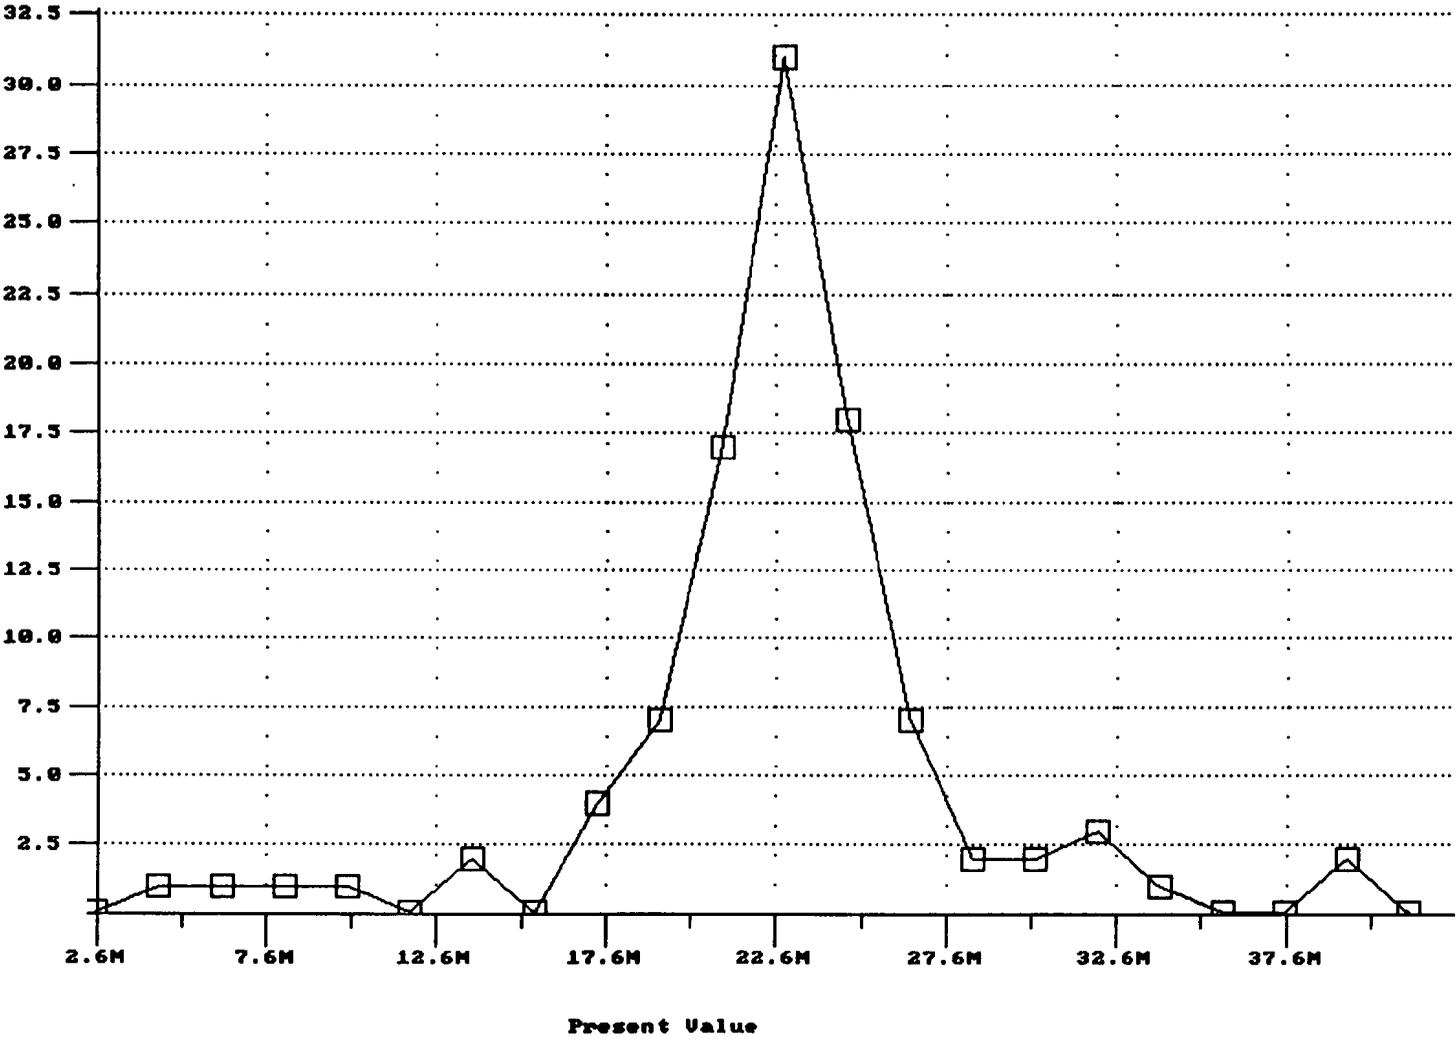
Optimizing interest crediting strategies is one of the uses of stochastic projections that returns the most bang for the bucks. If the mean increase and expected financial performance only goes up one basis point, mean annual expected earnings go up $100,000 for each billion dollars of account value. A five-basis-point improvement on $10 billion would give you expected mean earnings going up by $5 million annually.

While the potential rewards are great, so are the required stochastic projections. Fifty to one hundred stochastic runs for a crediting strategy would not be unusual. Again, a minimum of 100 scenarios is recommended.

Chart 16 shows the dynamically linked variables which need to be considered when testing alternative crediting strategies. Note that if you consider your dynamic lapse formula to be a SWAG (which is better than a WAG), that means it's not scientific, as opposed to something fixed and concrete, it would be prudent to triple the number of stochastic projections by utilizing your current dynamic lapse formula, one somewhat stronger and one somewhat weaker. In other words, you bracket your current formula. Note that a minimum spread can be very effective under a weak dynamic lapse rate formula, but a disaster under a strong dynamic lapse formula, where lapses are very sensitive.

CHART 16
Stochastic Interest Crediting

Competitor Rate

Dynamic Lapses
Formula
Parameters

Competitor
Recognition Tactics
Strategy
Minimum Competitor Rate
Weight

Financial Results

Minimum
Spread

Normal
Spread
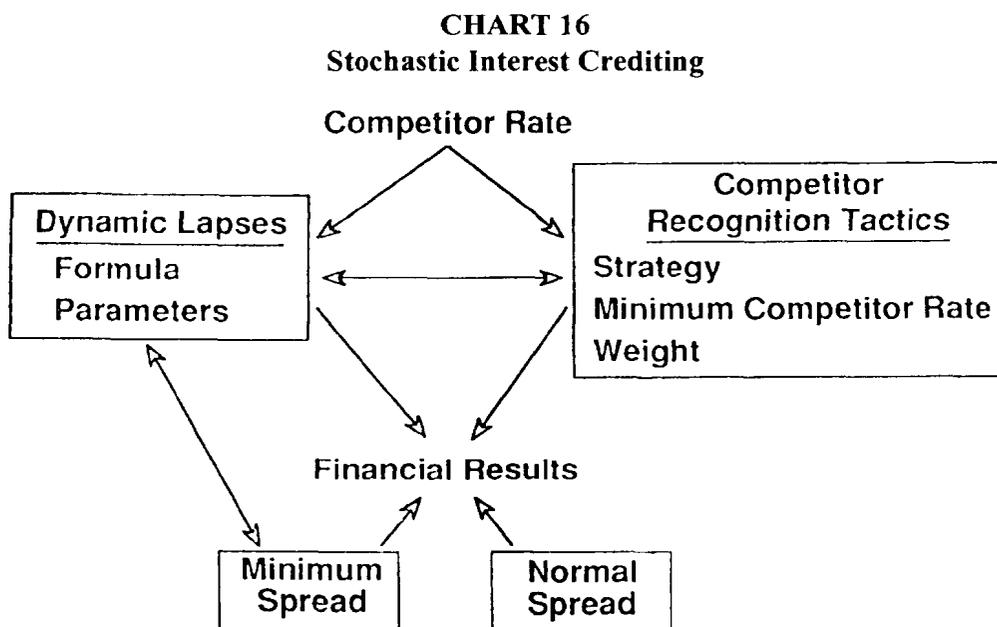
Chart 17 demonstrates a horizontal stochastic run which tests six interest crediting strategies at one time. The graph provides a risk/reward analysis based on ending surplus. You can compare the mean plus or minus one deviation and the best and the worst values, as well as the number of failures shown at the bottom. One of the biggest advantages of horizontal stochastic runs is that the output graphs provide you with a pattern of results that is most helpful when determining the next six or eight or ten strategies to be tested. Testing one strategy at a time is a slow and tedious process. You have to compare strategies to see how to improve them and select the next set of strategies to run.

Derivatives can be divided into four categories: forwards and futures, options on securities, interest index derivatives, and stock index derivatives. I won't go into much more detail on this.

The next topic is ALM and hedging strategies. How can hedging strategies help you? When do you buy derivatives? How much do you buy? How can you control the cost? Which derivatives should you buy? What hedging strategy is most cost-effective for a particular need? Stochastic analysis should provide answers to these questions today. In other words, what action or actions, if any, should be taken now with regards to hedging? Should you be hedging today? And, if so, how should this be accomplished?

Dynamic hedging strategies. In addition to showing you what to do now, a stochastic ALM system with dynamic hedging strategies will implement the selected hedging strategies throughout the projection period. It provides stochastic results with consistent risk profiles. Let's look at the various aspects of dynamic hedging strategies. The asset risk profile shows the vulnerability to the various types of interest risk. The risk profile is then used to develop dynamic risk factors, dynamic triggers, and dynamic notional amounts. These, along with dynamic cost controls, determine if you should buy derivatives, when you should buy them and how much you should buy in your ALM projection. We use the term *slice and dice*. This is used to fine-tune hedging strategies to arrive at an optimum balance between controlling the cost of the hedge and strengthening the risk profile. To put it another way, how can we get the most bang for the buck. In terms of stochastic analysis, this is a powerful technique for optimizing your hedging strategy.

CHART 17
Surplus
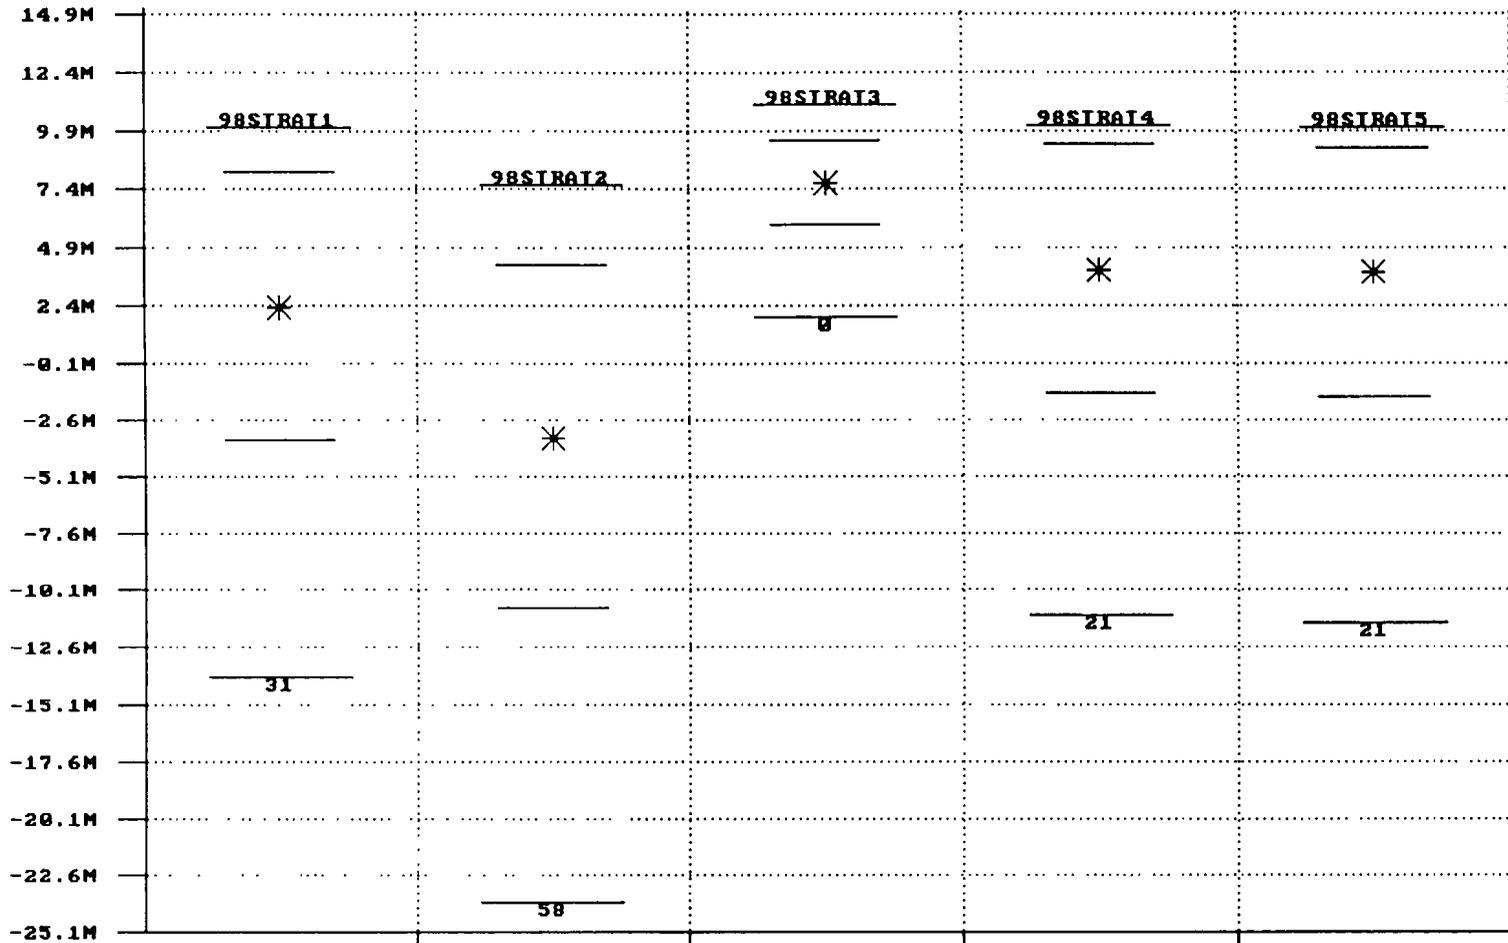
Number of Failures

—Maximum   —m+StdDev   ✳ Mean   —m-StdDev   —Minimum

The *slice and dice* technique is used for a family of two or more derivatives as opposed to a single derivative. The most common difference between these derivatives is the strike price. Here's an example. We just bought $10 million of five-year callable bonds yielding 8% for a new five-year 6.5% GIC giving us a spread of 150 basis points.

We have a risk of interest rates dropping, the bonds being called, and we have to reinvest at a much lower rate. This results in losing the 150-basis-point spread. Now the obvious way to protect your income from falling interest rates is to buy a hedge using interest-indexed derivatives. If the strike price is at the money, either the cost will be prohibitive or, if we have a cost limitation, very little coverage can be afforded. The answer is to buy out of the money. This is equivalent to co-insurance or a deductible. For example, we could buy a $5 million derivative, 75 basis points out of the money and another $5 million of nominal amount at 125 basis points out of the money. This means that if interest rates drop 4%, instead of having a negative spread of 250 basis points, we have a positive spread of 50 basis points when we reinvest. We are not fully protected. We suffer an affordable reduction in earnings instead of losing our shirts. The purpose of slicing and dicing is to come up with the best, most cost-efficient combination of derivatives. The most common variable between them is strike price. Does this take a lot of stochastic runs? You bet it does. Is it worth it? Considering the cost of derivatives and the available savings, you better believe it is.

Let's talk about an asset risk profile. There are four key components. When interest rates drop, how vulnerable are you to a dropping portfolio rate which will result in a reduced or negative spread, or for traditional life, reduced dividends or earnings? How vulnerable are you to a reduced spread because of the minimum credited rate in the contract? When interest rates rise, how vulnerable are you to higher dynamic lapses and/or reduced interest spreads? Also, how vulnerable are your market values? This is critical in terms of a meltdown point risk which we will discuss in a few minutes.

Dynamic hedging strategies can be used to protect: (1) protect your spread against the minimum credited rate, (2) investment income and your portfolio rate against falling rates, (3) income against rising interest rates, and (4) asset market values and stabilize GAAP earnings when interest rates rise. Another use of derivatives is to keep the duration of assets within specific boundaries.

The bottom line is derivatives can be very expensive. Investment bankers can't tell you the cost, but they cannot tell you if you should buy, when you should buy, how much you should buy and which strategy to employ. They also can't tell you how to allocate the cost between different strategies.

These answers require many stochastic runs. How much money and/or protection is being wasted because your company is not doing the needed ALM work to effectively utilize derivatives? Remember, the huge potential for increased protection, meaning a stronger risk profile and/or reduced costs, definitely justifies many stochastic runs with the accompanying analysis time. Derivatives can be very effective in the right situations if they are utilized wisely and cost-effectively. Ideally, derivative and investment strategies could be developed simultaneously. For example, callable bonds provide a higher return than noncallable bonds, but they have a risk. If the risk is too great, the derivatives could be used to reduce this risk, but this costs money and the cost may negate the additional yield to the callable bonds. The problem is to arrive at a set of strategies that provides an additional yield while reducing the consequences of these bonds being called to acceptable levels. In addition, the mean expected ROE or other measure of performance must exceed that of a strategy with no callable bonds or derivatives. Since there are probably several strategies that might accomplish this, the solution should be the best of these. Again, we are talking about a number of stochastic runs. This, of course, becomes practical only if the process can be performed in a reasonable time frame. If we're talking about the 10,000 scenario runs, one computer to do it overnight, it would take a little under six seconds per scenario. If it's a minute a scenario, then you figure ten computers; if it's ten minutes, you have 100 computers; if it's 100 minutes, forget it. I'm giving you rules of thumb for how long and how many computers you need.

Chart 17 shows a graph produced by a horizontal stochastic run comparing five hedging strategies using risk/reward criteria.

Let's move on to the most timely and critical topic of this presentation—real world risk analysis. Real world risk analysis means we fully implement the first step of an effective ALM process per the July 23 memo from the ALM Principles Task Force, which is to identify the level of risk exposure. We are doing this for interest rate risk. We are not doing it for economic risk. Both the

industry and the ALM actuaries have been very lucky the last 12 or so years. The economy has been stable, but, for the most part, improving. This has been in addition to stable, moderate-to-low interest rate levels. With the current world economic situation, there's a reasonable possibility that our luck may be running out. Why do I say this? Because of the state-of-the-economy in South Korea, Thailand, Malaysia, Russia, and of course, Japan. Here are some quotes from articles in the September 5 issue of *The Economist*. One is titled, "Heading For Meltdown?" "For the first time since the early 1980s, global slump as a thinkable, even plausible outcome. Indeed, in some ways the danger now is greater than it was then." "What mainly stands between the world and an economic setback worse than anything since the Great Depression of the 1930s is the present momentum of growth in America..." In another article I read, "At the annual meeting of the Federal Reserve Bank" . . . "some central bankers were privately admitting that these are the worst global economic conditions they have seen in their lifetime."

Does all this mean we are going into a depression or even a major recession? No, not necessarily, But it does mean that these are now distinct possibilities. What this could mean to your company can no longer be ignored with impunity. The current worldwide economic conditions make it imperative that we identify the level of risk exposure, and then take the appropriate actions to deal with any levels of unacceptable risk.

Let's talk about the meltdown point. Meltdown points exist in the real world, but they are ignored in our ALM projections. A meltdown point occurs when a published ratio sinks to a specific "unacceptable" level. In the real world when this happens, an investigative/financial reporter will write an expose that company XYZ is in trouble. This article gets reprinted in other publications, and other reporters publish new articles. Policyholders get nervous and lapse their policies. This leads to new articles, and this vicious cycle leads to a run on the bank, ala Mutual Benefit and First Capital. Regulatory projections are only for insurance in-force and assume starting reserves equal to beginning assets. Real-world projections include new sales with a positive beginning surplus—the more the better. The greater a company's surplus, the less vulnerable it is to adverse conditions. Thus, beginning surplus matters.

Finally, a meltdown point occurs when key published ratios drop to a specific level. The key ratios include the ratio of the market values to the book value of assets, and the critical risk-based capital ratio. Note that these are total company ratios. This means your stochastic projections must be done for the total company to evaluate this risk. What happens to the scenario when a meltdown point is hit? That company goes into receivership. There are no more dividends to stockholders and there is no more ROE; in other words, it is a disastrous scenario. How does the state-of-the-economy affect assets? A bad economy has a major impact on assets. Defaults and foreclosures jump, market values drop, for some companies significantly, for two reasons. One, in a bad economy, the spreads over Treasury widen. Thus, if you buy BB bonds in a good economy and then a major recession hits, the large increase in the BB spread over Treasuries will zap your BB market values. The second reason is that there will be a significant increase in rating migrations. In other words, more AAA bonds will turn into AA bonds and so forth. Interest rate scenarios alone do not recognize these risks. Last, certain types of high return assets lose their liquidity in a bad economy.

A bad economy produces major dynamic affects on liabilities. Policy loans jump, not because your credited rates or dividends are not competitive, but because people who have lost their jobs, are working at lower paying jobs. They have reduced income from investments and need money to pay their bills. This is the same reason that lapses also dynamically jump, and why premium persistency suffers. Depending on the market, there could also be an effect on the level of sales. Notice that all of these act to impact insurance cash flow and force asset disinvestment while asset market values are being forced down.

Averages do not show the risk. The average default rate over the last 10 or 20 or 50 years are meaningless in terms of risk. The risk is shown by peaks that occur in a bad economy. This also holds true for foreclosures, policy loans, lapse rates, premium persistency, and so on. The interest-scenario-only stochastic runs ignore this risk. Remember, increased return means increased risk. The increase in the average default rates do not show the total risk in investing in BB versus AAA. Our interest only stochastic projections produce an increased return without showing the full accompanying risk. Again, we need to identify the level of risk exposure.

Table 3 has 100 interest scenarios stochastically run. We are comparing strategy A and strategy B. Strategy A has a more aggressive investment strategy, although it could be crediting strategy. Strategy A produces a higher mean ROE with a higher probability of meeting the company's financial goals, and it has a higher probability of meeting the company's threshold for minimal financial goals, let's say a 6% ROE. Based on our interest-only stochastic runs, strategy A sure looks much better than strategy B. There's no way we would be misleading ourselves is there?

Table 4 shows a 10,000 scenario run that compares strategies A and B. Strategy A still produces the higher ROE. It still has a slightly higher probability of meeting the company's financial goals, and has an equal probability of meeting the threshold goal, but we're now paying a price. The probability of having our ratings drop is three times as high, and the probability of hitting a meltdown point is five times as high. Would you be willing to pay the price of this additional risk for an additional 1% of ROE? Would your management do this? Would Mutual Benefit have made the same decisions if they had access to this type of information? Would First Capital have done this?

We have seen that you cannot identify the total level of risk exposure with interest-rate-only stochastic runs. So what are our choices? We can bury our heads in the sand and assume it will never happen here and the economy will always be good. Or we can take the next step and again identify the level of risk exposure. In other words, start performing total company interest rate, state-of-the-economy, vertical stochastic runs. To do this we need an action plan—a process for doing this. First, we run a total company 10,000 scenario stochastic projection. Do we have problems? Is the meltdown point risk too high? Is the ratings drop risk too high? If so, the next step is to run horizontal 10,000 scenario projections for each major strategic business unit and select better, safer strategies. Then rerun the total company vertical stochastic projection and analyze the results. If it's still not good enough, you repeat the process until it is. The end result is discovering and then managing your risk.

MR. JACKSON: I do have one question myself. Mel, you did give us some run time quotes, Mel, you didn't specify the system you referred to.

TABLE 3
100 Interest Scenarios

	Strategy A	Strategy B
Mean ROE	15%	13.5%
Goal Probability	80	76
Threshold Probability	95	90

TABLE 4
10,000 Interest/Economy Scenarios

	Strategy A	Strategy B
Mean ROE	14.2%	13.2%
Goal Probability	70	68
Threshold Probability	82	82
Ratings Drop Probability	12	4
Melt Down Probability	5	1

MR. STEIN: I was giving the required run times per scenario for a single PC to run 10,000 scenarios in a single night.

MR. JACKSON: There was reference to your system. Was it CHIMP?

MR. STEIN: Not really. I was just trying to give a rule of thumb for any system. In other words, if you want to get it done on one computer overnight, you have to use a six-second scenario. The reason I'm not more specific is that your run may start at 4:30 p.m. or 5:00 p.m., and you have to end at 8:00 a.m. or 9:00 a.m. So whether it's 15 hours or 16 hours or 14.5 hours, it will be 5.7 or 5.8 or 5.9 seconds. CHIMP can do that sometimes, but not always.

MS. ANDREA MAXINE THOMAS CHAPMAN: I want to ask a question relative to my expectations for this session. I saw mention of programming languages and systems requirements. Are you saying that insurance companies need to purchase software? This is not something that you want to begin to develop within your company. We've developed a home grown kind of model, and we do about ten scenarios or liabilities. The assets are pretty straightforward. We're purchasing software because the regulators are beginning to ask us for more and more scenarios, and all I'm hearing here is you need to purchase them; I'm not hearing how I can enhance what I have in-house. Is the assumption that I'm making true?

MR. WHITE: Actually, all of the software languages that we refer to are commercially available. This, as in any modeling, requires regular time, regular trial and error, and some companies choose to outsource it and some companies do choose to develop internally. There is expertise to be gained internally as well as externally, and that's something that you'd have to balance out for yourself. In our case, it's our intention to suggest that the computer power exists, the languages are there, and it can be done. If you would choose to outsource it, we'd love to talk to you. You certainly could do it.

MR. JACKSON: Let me respond. I think Mel Stein stuck a little bit too close to not being commercial about his software, PTI software. It does take feeds from some of the major systems as does Larry's, and it does seek to provide and use those feeds to produce some more efficient, quicker response. He was sticking a little bit closer to the letter of the law. He does have software available for sale that would improve stochastic run times.

MR. LEE R. LAMBERT: Can you give us any guidance about reducing the number of cells, and getting the task down to one that can be done in a few seconds? I have a very large model that I'm running on PTS and if I do it completely, it takes me about ten hours so I can't do very many scenarios. How many cells would you typically have, and what guidance can you give us in making that more manageable?

MR. JACKSON: I'll take a shot at that first because I have used PTS considerably. I think I've run into situations with client companies that have used PTS, and I know of one situation where some of the modeling of the cells represented up to 30% of the liability cells and characteristics. Typically, it's rate for a system that runs efficiently in PTS to have more than 5% of the cells represented of your total in-force business. A 0.5% or 1.0% for PTS is a more realistic way to go. With PTS, one of the questions is whether you're interfacing with the Intex system which takes significant time to run to get the accuracy. These are some of the challenges I know that do exist with PTS. That is a pretty robust system, but one of the challenges with that system is time. One of the ways you can deal with it is with those kinds of cutting down of your liability modeling or considering tying in to other systems to help you run efficiently.

MR. WHITE: I have a comment. I think you hit the nail on the head there. The first thing is get a very fast computer. One key is as you take information away from your model, you make your model less effective. So you have to be judicious about what you take out and what you leave in. In our case, we typically choose to take a broad representation of the securities that are available in the market, and yet cut down on some of the scenarios. The theory behind that is reasonably well supported. We also found this information empirically, such that, for any given business line or for any given company, there are certain scenarios that are constraining or are what we call boundary type scenarios. The rest of the scenarios typically are noise. If you can effectively identify which scenarios are the most difficult for the asset managers to solve, then those are the ones you throw into the solver. Give the solver the most difficult problem. We typically find about 95% robustness against a smaller universe of scenarios, once we've identified the scenarios that are the ugliest ones.

MR. STEIN: I'll add my two cents worth. All the major vendors, PTS, TAS, or whatever, have these model generators. You can crunch down the number of cells and pretty well match the first year's results. The problem is, after the first year, they start to diverge. It has been our experience that there are a number of companies that we're waiting on because they cut their models too much and now they're trying to build them back up again to get better results they can rely on more. They

are sort of between a rock and a hard place. With CHIMP, the original number of cells for traditional life, for example, is immaterial. We're about to do one now that has a high six-figure number for the numbers of cells. We're going to cut the run time from many hours on a mainframe to well under a minute on a fast PC.

I'll let my esteemed colleague talk about the assets. On the liability side, it's not that easy to chop down the number of cells in a model and still get the long-term accuracy. In the first year, you can always force it, but the results may diverge after that. You have to run it both ways to really see how accurate your reduced model is.

MR. JACKSON: I'm not sure that completely answers your question. I know it is a problem, and I've run into the same kinds of challenges. That's kind of why we put this session together. We want to see what kind of solutions the panelists have and if some people out there do have some suggestions. How do they address the challenge that all of us have seen?