Innovations in Investment Management

Cutting-Edge Research from the Exclusive JOIM Conference Series

EDITED BY
H. Gifford Fong

FOREWORD BY Martin L. Leibowitz, PhD

BLOOMBERG PRESS
NEW YORK

CHAPTER 1

A New Generation of Pension Fund Management

ROBERT C. MERTON

Who—Robert C. Merton received the 1997 Nobel Prize in Economics for a new method for pricing derivative securities. He is the John and Natty McArthur University Professor at the Harvard Business School, past president of the American Finance Association, and a member of the National Academy of Sciences. In 1993, he received the inaugural Financial Engineer of the Year award from the International Association of Financial Engineers. Formerly a senior adviser to Salomon Inc. and JPMorgan and cofounder and principal of Long-Term Capital Management, Merton is currently chief science officer of Trinsum Group, Inc., which offers SmartNest, a pension-solution system. Merton has a BS in engineering mathematics from Columbia University, an MS in applied mathematics from the California Institute of Technology, and a PhD in economics from the Massachusetts Institute of Technology.

New Thinking—The legacy of the defined-benefit (DB) plan might be the defined-contribution (DC) plan, but the legacy of the DC plan will replace both the DB and the DC plans. Although still in the early stages of implementation, this "next-generation" plan will be able to provide the solution of predictable cash flow for life, without credit risk, offering the cash flow profile of a DB plan, but through a model derived from the subscriber's DC assets.

Introduction

In talking about pension plans at this point in American economic and corporate history, we need to discuss three linked issues: the DB corporate plans that worked for our parents; the DC plans we're getting today,

because corporations no longer want to bear the expense and risk of the DB plan; and how we can reshape the DC plan into something that feels more like a DB plan in the near future. Although this chapter focuses on a U.S. context and corporate pension plans, retirement and asset-management issues are a global challenge, and, indeed, while the details vary across geopolitical borders, the fundamentals are shared by employers and employees across the globe. What's happening with the DB plan, why is it happening, and what are the implications for asset markets and asset management?

I. The Defined-Benefit Plan

In 1990, DB plans were approximately 22 percent of the retirement market in the United States. The DC percentage was pretty close, but a little bit smaller. By 2004, however, DB plans represented only 14 percent, and DC plans were almost double that. In the year from April 2005 to April 2006, of the 627 Fortune 1000 companies intending to freeze or terminate their DB plans, 113 had already essentially either ended or frozen them, and the announcements of DB plan terminations were up 60 percent from a year earlier.

The watershed event occurred in January of 2006, when IBM, long known for being employee-centric, highly rated, and overfunded, decided to end its DB plan—even for existing employees. There have been several other major defections since, the most recent being DuPont. What's happening here?

Well, it's safe to say that companies are *not* discontinuing their plans because their employees marched on their HR departments to say, "We hate defined-benefit plans." On the contrary, employees loved the plans. The problem was the plans were too costly, both to fund and to carry the risk on the balance sheet. But this awareness of the cost factor was not entirely unanticipated. Some academics and financial practitioners were well aware of it long ago, and they even published analyses explaining why the existing accounting treatment would lead to the underestimation of plan costs.

How Did the DB Plans End Up Being So Costly?

Pension accounting was almost surely the key reason that the costs of DB plans were systematically underestimated. Under the accounting rules, the plan sponsor's projected future contributions to support benefits were determined by assuming that future expected returns on the pension fund assets would be realized, without making any adjustment or provision for the risk of those assets. Pension advisers thus saw that pension funds, by holding higher-risk assets with higher expected returns, would lower the calculated future contributions by the plan sponsor. Since future

contributions are the cost of the plan, the advisers concluded that, by investing more heavily in risky assets, principally equities, the cost of the plan could be reduced.[1] This practice, which makes no adjustment for risk, in effect treats the expected returns on the portfolio and projected future contributions as "sure things."[2] This pension accounting treatment is equivalent to applying an above-market, risk-free discount rate to compute the present value of future pension benefits, which will understate the cost estimate for plan liabilities.

If, for example, the risk-free rate is 5 percent and the expected return on the portfolio is 9 percent, there's a 400-basis-point spread. If you are calculating pension liabilities with an effective duration of ten years and you used 9 percent instead of 5 percent, you'd get two-thirds of the cost. In other words, if you discounted at 9 percent over ten years instead of 5 percent, you'd end up with two-thirds of the true present value of the costs. Now try flipping that over: If you used the 9 percent discount rate and forecasted your needed contributions based on that, in effect you'd be treating the liabilities—or the cost of servicing those liabilities—as if they were two-thirds of the actual cost. What you thought was a dollar's worth of benefits was really costing you a dollar and a half.[3]

It wasn't that the unions or the employees were any smarter than the corporations were in understanding that the benefits were really worth more than they appeared to be. It's just that, generally, when someone offers you a choice of two things, and one of them is considerably more costly to produce, you usually perceive the value and choose the costlier one, even if you don't explicitly know its price. If my company was giving away company cars and I offered you a choice of a Toyota Camry or a Bentley, and you had no clue as to their prices and could never sell them, I'd still bet that you, along with most people, would choose the Bentley—just because a $200,000 car is likely to be more attractive and have better features than a $30,000 car.

The second reason that the DB plans ended up being so costly was that when the volatility of the pension assets' earnings and contributions was smoothed, the results were both a higher projected rate of return and an unrealistically lower reported volatility rate, which led to an inadequate provision or projected contribution reserve for the risk of the assets. But all of this was already understood, at least by academics.

When, in the 1990s, the markets in the U.S. and the U.K. were going up, it wasn't a problem that caught plan sponsors' attention. Between 2000 and 2002, however, a convergence of circumstances led to a substantial drop in global stock markets and falling interest rates, which resulted in falling pension assets and rising pension liabilities, that weakened company balance sheets and caused large companies in already troubled

industries to go bankrupt. All of this surely *did* capture plan sponsors' attention. This is what really killed the DB plans, and this is why they're not coming back. They cost a lot more than plan sponsors *thought* they cost.

Then, at the same time that the inhospitable economic climate was putting corporations under financial stress, something else was happening: Changes were proposed to accounting standards[4] to make them more accurately reflect valuations and volatilities. The eventual implementation of those standards will have a material impact on companies' balance sheets and on earnings volatility—because the companies are holding equities in the pension funds—that will extend far beyond the economic cycle.

All of that said, there is still more than $2 trillion in pension assets left in corporate DB plans in the U.S. alone, which are going to be there for a long time. Also worth noting is that as pension accounting changes to reflect economics and risk more realistically, it is likely to create an investment trend within the DB area: I predict that existing DB plans will move in the direction of matching the duration of the liabilities with the duration of the assets—plus alpha. I'm not predicting total immunization, but, rather, *moving toward* matching the duration—plus alpha.

The Cost of Holding Passive Risk

Market risk, systematic, indexed risk—not alpha, not superior performance—is clearly a passive risk. Anybody can get that, and it's a zero net present value (NPV) investment. A firm can't buy a million dollars' worth of the Standard & Poor's 500 and expect the market to value it at more than a million dollars. Is there a problem with corporations holding passive investment like that in their pension plans?

Well, yes, there is a definite cost to corporations for holding passive risk:

- ***Risk identity.*** The balance sheet of the firm[5] (not just the accounting identity, but also the value identity that assumes that the value of the left side has to equal the value of the right side) also has a *risk* identity. That risk identity holds that the total volatility of the left side has to be borne by the right side—by debt and equity. If you accept that, then, once you determine a desired credit rating, which specifies the amount of risk that the debt will bear, there is no flexibility regarding the amount of equity you need. Given the riskiness of the assets, the amount of equity you need is completely determined by the amount of debt you have and its desired credit rating. That's the risk identity, and in any consideration of total risk, volatility of assets drives the amount of equity you need.
- ***The frictional costs of volatility.*** Even if you hold a passive asset like the S&P 500, it still creates volatility in the asset base of the

contributions are the cost of the plan, the advisers concluded that, by investing more heavily in risky assets, principally equities, the cost of the plan could be reduced.[1] This practice, which makes no adjustment for risk, in effect treats the expected returns on the portfolio and projected future contributions as "sure things."[2] This pension accounting treatment is equivalent to applying an above-market, risk-free discount rate to compute the present value of future pension benefits, which will understate the cost estimate for plan liabilities.

If, for example, the risk-free rate is 5 percent and the expected return on the portfolio is 9 percent, there's a 400-basis-point spread. If you are calculating pension liabilities with an effective duration of ten years and you used 9 percent instead of 5 percent, you'd get two-thirds of the cost. In other words, if you discounted at 9 percent over ten years instead of 5 percent, you'd end up with two-thirds of the true present value of the costs. Now try flipping that over: If you used the 9 percent discount rate and forecasted your needed contributions based on that, in effect you'd be treating the liabilities—or the cost of servicing those liabilities—as if they were two-thirds of the actual cost. What you thought was a dollar's worth of benefits was really costing you a dollar and a half.[3]

It wasn't that the unions or the employees were any smarter than the corporations were in understanding that the benefits were really worth more than they appeared to be. It's just that, generally, when someone offers you a choice of two things, and one of them is considerably more costly to produce, you usually perceive the value and choose the costlier one, even if you don't explicitly know its price. If my company was giving away company cars and I offered you a choice of a Toyota Camry or a Bentley, and you had no clue as to their prices and could never sell them, I'd still bet that you, along with most people, would choose the Bentley—just because a $200,000 car is likely to be more attractive and have better features than a $30,000 car.

The second reason that the DB plans ended up being so costly was that when the volatility of the pension assets' earnings and contributions was smoothed, the results were both a higher projected rate of return and an unrealistically lower reported volatility rate, which led to an inadequate provision or projected contribution reserve for the risk of the assets. But all of this was already understood, at least by academics.

When, in the 1990s, the markets in the U.S. and the U.K. were going up, it wasn't a problem that caught plan sponsors' attention. Between 2000 and 2002, however, a convergence of circumstances led to a substantial drop in global stock markets and falling interest rates, which resulted in falling pension assets and rising pension liabilities, that weakened company balance sheets and caused large companies in already troubled

> firm, and thus some amount of equity has to be held to cushion that asset's volatility. Since equity has significant frictional costs, and that asset has a zero NPV, holding a passive asset is a losing proposition. Even if the asset is fairly valued, to hold it you have to have additional equity to cushion its volatility, and since equity has frictional costs, it is a net-negative deal. As a practical matter, it can be very hard to strip out some of those passive risks at an acceptable price, except in the area of pension funds, where it is relatively easy to eliminate one type of risk: systematic market risk.[6]

The frictional risks of holding passive investments in the firm are the tax disadvantage of equity versus debt financing (debt interest payments are tax-deductible and equity dividends are not) and the discount in public equity prices from the inherent conflict of interests between owners and managers of the firm, which we'll call "agency cost." We don't *have* a precise number to assign to agency cost, but we know it's significant.

If you concede that there are frictional costs to holding passive assets in the firm, then you have to concede that it's a losing proposition, and, if it's a losing proposition, then you have to consider that the risk profiles of the pension funds affect more than the pension plan beneficiaries. Rating agencies and investors are coming to recognize that pension plans need to be integrated into corporate risk assessments and that risk mismatches between plan assets and liabilities are large. The agencies have already been looking at the deficits and surpluses, but now they're also starting to think about risk mismatch between assets and liabilities. In addition, the FASB changes previously mentioned, when they become effective, will make more transparent these risk mismatches through greater reported earnings volatility. These are the reasons why you'll see corporations moving toward not taking systemic risks.

Alpha risks can be justified, and alpha-generating investments will continue to be included, because, if you have a positive alpha after fees are deducted, that's equivalent to a positive net-present-value project, and then the scale of the project is determined by the trade-off between the value increment from the project and the added frictional cost of equity to support its risk.

If my prediction were to prove correct, that existing DB plans will move in the direction of matching the durations of liabilities and assets, what would be the implications for the capital markets? What if all the pension funds sold off all their passive equities exposures and bought debt (while, of course, retaining any alpha generators)? What would that do to the markets? Would it crush the equity markets and cause debt prices to rise and interest-rate spreads to tighten?

It's not likely to, at least in any sustained sense. Just imagine, for example, if General Electric Co. (GE) and its pension plan sold IBM stock and bought IBM bonds, and IBM sold GE stock and bought GE bonds. Imagine that all the pension funds sold their stock and bought bonds. What would happen then is that, because the risk of the left side of the balance sheet was going down—because the pension funds would be getting out of equities and into debt—firms wouldn't need as much equity to maintain the same credit rating. They would be permitted by the market to recapitalize their capital structure as a consequence of lowering their asset risk; they would issue more debt and buy back their own equity. They wouldn't need as much equity as a cushion, because they would have gotten rid of passive risks. If you follow that line of thinking, it would turn out that, when GE was selling IBM stock, IBM would want to buy its stock back and issue bonds. That would mean that the bond spread that you thought was going to tighten from greater fund demand wouldn't necessarily tighten at all, because bonds would be being issued to provide extra supply.

And then guess what IBM would be doing: IBM would be selling its GE stock, and GE would be buying its stock back and issuing bonds to finance the purchases. So there would be no sustained disequilibrium; the whole process would not, after all, destroy the equity market.

What you would observe after this is that debt-to-equity ratios would go up. If you examine U.S. pension plans, you'd see that they look very much like Japanese cross-holdings: GE is holding IBM stock in its portfolio, and IBM is holding GE stock in its portfolios, and there's an awful lot of stock there that's going nowhere.

II. The Changing Asset-Management Process

Taking It Apart: The Bifurcation into Alpha and Beta

What are the trends we're seeing now in asset management, and what do those trends imply for the future? The biggest one is something you could call "barbell investing."[7]

There has been huge growth in asset management businesses offering pure exposures to beta, such as index funds, exchanged-traded funds (ETFs), and derivatives, which are very efficient, low in cost, and the most scalable.

The other area where we've seen tremendous growth, clearly, is in "pure" alpha generation, from hedge funds, private equity, and direct real estate funds. The growth has been at the two ends, pure beta and more-or-less pure alpha generation.

The cost of the barbell growth has been to the middle, the traditional style of asset management, which mixes the two together and is essentially a long-only style. (Many of the star managers in the middle, in fact, have left to start their own hedge funds.) What are the implications of this barbell investing?

One implication is that institutions will move into alternatives, which we're seeing, whether outsourced or captive. Now think about the way this will work, beginning with a status quo equilibrium analysis. Even if institutions end up with holdings identical to what they had before this change in institutional structure, it should still *appear* that they are going into alternatives. Why is that?

The answer is bifurcation. What used to be all in the middle is now bifurcated. Institutional investors must move to new places to re-create what they had before—perhaps more efficiently, perhaps not. Thus, the trend toward alternatives is almost going to have to happen just because of the change in the asset-management-business structure, whether or not there are any extra returns to be generated from it.

The second implication of barbell investing regards the transparency of strategies and the source of value added. You pay one fee rate for alpha and quite a lower one for beta. When the two were plugged together, that wasn't very transparent. But this new structure is going to be much more transparent, and the composition of each component of the whole will be much more easily measured. The very existence of the bifurcation, the separation of alpha and beta, is providing a transparency that is making pension funds much more sensitive and fund managers more aware of what they're doing. This, in turn, will have the effect of accelerating the dynamic: Alphas will become betas as new exposure becomes systematized and resulting fees go down for a given strategy. After all, what do alpha and beta mean? If an investor is doing the best he or she can, and the investor's manager can improve the investor's risk-return frontier, that manager has created an alpha for that investor, however he's done it. It might not be an alpha for a smarter investor who has already incorporated it, but everything is relative.

To illustrate what I mean when I say alphas are becoming betas, consider the small-stock-versus-large-stock risk premium, or consider the value-versus-growth risk premium. If your investor client hasn't been using either of those strategies before, or if his benchmark had been simply the broad market undifferentiated, then you could create an alpha for that investor simply by tilting into small stocks or tilting into value stocks. You would be creating value for your investor because you would be doing better for him or her than he was doing before. Once the source of added return is recognized, however, then very efficient means

arise for creating those differentiated exposures, and, so, from a fee point of view, the strategy would be taken out of the alpha class and put into the beta class.

You would see similar effects from momentum (if it exists) or from any other strategy that, once recognized and systematized, could be provided at low cost. Once that happens, why would anyone pay a premium fee for the asset management? That's what I mean when I say that by bifurcating the industry the dynamic process of alpha generating strategies transforming into beta strategies will accelerate—for better or for worse.

Putting It Together Again: The Assembler Function

So, now that we've separated the alpha and the beta, what's next is putting them back together again—assembling them into the institutional portfolio, but in a totally different way from the preset, blended style we had before. Institutions are going to be looking for skilled assemblers who can take the two ends of the barbell, the alpha and the beta, and assemble them in a much richer, more sophisticated way.

The function of the assembler could be to take pure alphas and mix them with pure betas, or perhaps create some sort of transparent prepackaged mix of specific beta exposures plus alpha, like the S&P 500 plus alpha, or the Morgan Stanley Capital International funds (MSCI) plus alpha, or whatever beta index you like plus alpha.

I believe there are already investment firms doing this—not offering pure alpha, but offering a combination: having the pension-fund investment committee choose the index to combine with the alpha. Institutions will need to decide on the mix of indexes to use as well as the mix of alpha providers. The assemblers to perform this role can be inside the institution or outsourced. The need for people in the assembler function is going to grow tremendously and have a significant impact on the industry.

What skills will an assembler have to have? He'll have to develop new, integrated, and consistent methods of measuring risk for alternative and traditional asset classes. Whether you're talking about hedge funds, private equity, or real estate, if you're going to be assembling alphas and betas in pension-fund portfolios that include them, you will need to have consistent risk measures to be able to combine them with traditional asset classes of stocks and bonds.

Let me give you an example: When you're dealing with alternatives that don't have much of a secondary market, such as private equity or direct real estate investment or structured products such as collateralized debt obligations (CDOs), you're going to have to take into account the effect of stale pricing. That's a problem we've always had to deal with, and it principally affects the correlations between returns in measuring the

contribution to portfolio risk. But let me tell you about an experiment that I ran that illustrates the impact of stale pricing.

I created a firm, a hypothetical private-equity fund, called "S&P 500 Private." Now, just between us (don't tell anybody), the only asset of that portfolio was the S&P 500. My hypothetical portfolio provided its investors with a weekly, not a monthly, NAV.[8] Although the NAV was reported each week, the portfolio was priced only every two weeks. I knew the true price of my portfolio, the S&P 500 Private, today, because I had just priced it in the market, but I knew that next week I wasn't going to reprice it. The reported price next week thus would be a stale price and identical to the reported price this week. Then, the following week, I'd take into account whatever happened for the two-week interval and reprice the portfolio to market.

Now you have a picture of the portfolio's reported performance: Since it was being repriced to market every other week, it had a very systematic pattern. It would look as if there had been no price change for almost two weeks; then it would move either up or down; then no price change for almost two weeks; then up or down. The price dynamics were very transparent, with stale prices reported every other week.

Now let's assume that we wanted to measure the performance of S&P 500 Private, relative to our benchmark, the S&P 500. What did we know was the true alpha and beta of S&P 500 Private? Its beta was 1.00 and its alpha was 0.00, because the benchmark and the portfolio were identical.[9] What if I had run the weekly reported NAV against the benchmark? The average return, or average excess return, would, of course, have been unaffected, because the average return of any portfolio depends only on the beginning and ending points of the sample period, whether you sample once a week or once every two weeks.

But how would the estimate of the beta be affected by stale pricing? We know the true beta is 1.0, and we know the beta equals the ratio of the standard deviations of the fund's and the benchmark's returns multiplied by their correlation.[10]

What got hammered by this systematic stale pricing, however, was the estimate of correlation between the portfolio's and benchmark's returns. In reality, the actual correlation was perfect; the true portfolio value and the benchmark moved exactly in tandem. But, because we were using stale prices half the time—regardless of how the S&P 500 benchmark was actually performing—the reported NAV didn't move in half of the weeks. This produced a measured correlation with the benchmark a little under 0.5. Estimated beta, therefore, instead of being 1.0, was 0.5. We had understated the systematic beta risk of the fund.

Now consider the estimate of alpha, which equals the average excess return of the funds, less the beta, times the benchmark's excess return.

Since the average returns didn't change, and the estimated beta was understated by half, the alpha, by definition, has to be equal to half the realized excess return on the benchmark—in this case, the S&P 500. Guess what: When I ran this for the five-year period from 1995 to 2000, the annual estimated alpha came out to be 1,000 basis points!

It sounds like a joke, but that's how it works. Granted, 1995 to 2000 was a pretty hot period, with the excess return on the S&P 500 roughly 20 percent, and because the estimated beta was understated by half, the estimated alpha was roughly half of 20 percent, or 1,000 basis points.

This is a fast pass on the effects of stale pricing on performance measurement, but you can see intuitively—the point of this transparent exercise—that it's an overestimate of alpha caused by a very simple temporal misalignment of returns. The point is that these errors in the estimate of alpha can be very, very large, and, therefore, if you don't correct for stale pricing, it's not just a round-off error; it can have a huge impact on estimated performance of a fund.

With an actual private-equity portfolio, the stale-pricing problem is far more severe than in the hypothesized example: It doesn't get marked to market every other week; it gets marked only when there's a new round of financing, or an initial public offering (IPO), or another exit, or, occasionally, a write-off. So it could be a rainy Thursday night when one of those events occurs. Now you have some idea of what that does to the estimated correlation. This is not to say that private equity as an asset class has not earned superior excess returns. It does say that, in the world we're in, when you get into this assembly process, this integration of the pension-fund portfolio components, it's going to be very, very important to take all of this into account. Tools for doing so exist, but the way we measure performance will have to be very different from what it was in the past.

Sustainable Returns for Alternative Investments

Let's talk about sustainable returns for alternatives, in particular, for hedge funds. The question is, What are the sources of sustainable returns for hedge funds? Let me offer three:

*1. **Market exposure.*** Even managers of hedge funds that try not to, can, in very subtle ways, have systematic stock-market exposure from their trading strategies, such as in convertible-bond arbitrage. They may think they have stock-market risk hedged out if they short the underlying stock against a long position in the convertible, but, if they don't hedge out the volatility exposure of the convertible, and they're exposed to changes in volatility, then they have implicitly taken risk in the stock market even without intending to, because stock returns and changes in volatility are negatively correlated.

2. Liquidity-event risk. What I mean by *liquidity-event* risk is not bid-/ask- spread liquidity, but the kind of liquidity event that shocks the market, causing market transactions to shut down. In such a market, until market players can figure out what's going on, it will stay that way, and there's a flight to quality.

My colleague George Chacko studied returns on the CSFB/Tremont Hedge Fund Indexes for all styles of hedge funds from 1994 to 2004. What he found was that, by using a single factor from the equity market for liquidity—even though hedge funds of different styles take positions across all the markets—much of the residual alpha performance for all styles became insignificant.

Consider corporate bonds in terms of liquidity exposure. For exchange-listed equities, the median time between trades is perhaps a couple of minutes, depending on the float. For corporate bonds, however, the median time between trades is six months, because most corporate bonds don't trade at all. There's a tremendous difference in liquidity between corporate bonds and equities.

Again, consider convertible-bond arbitrage. You go long the (illiquid) corporate bond, and you short the (liquid) equity. So you do have liquidity exposure, whether you want to or not. Or look at long-short equity arbitrage. Typically, the short equity position is in the larger-cap stocks that are more liquid, easier to borrow to short sell, and less likely to get squeezed out of the short position,[11] and the long position is typically in smaller, less-liquid stocks.

There are many sophisticated strategies being used, and there are delays and lags, but it appears—at least from the data we see—that hedge funds are systematically taking on liquidity-event risk. There's nothing wrong with that, but investors in hedge funds should be aware of it.

3. Institutional rigidities. From a financial-system perspective, I think this risk is the most interesting, and it's one I've believed in for a very long time. What is it that all hedge funds have in common? Certainly not strategies; they undertake all different strategies. The one essential characteristic that makes a hedge fund a hedge fund, the commonality that binds them all together, is that they're minimally regulated institutions, certainly relative to mutual funds, banks, insurance companies, or even pension funds, or any other kind of funds.

Institutional rigidities can occur from many different factors—accounting rules, regulatory agencies, taxes, asset classes, geography, transaction times, or methodologies, to name a few. What I mean by the impact of rigidity is the situation that occurs when an investment manager would *like* to take a position, or get rid of a position, *but is prevented from doing so by some controlling factor outside the trading strategy.* There are many such

rigidities, and the changes that come with globalization mean that the web of rigidities is becoming ever more complex, because of regional differences in regulations and standards.

The markets that are perhaps most affected by institutional rigidities are the credit markets, especially in the area of accounting. In the equity markets, almost everyone's accounting is already performed on a mark-to-market basis. In the credit markets, however, while some players are mark to market, a vast number of them are not. That accounting difference gives rise to different tactics and accommodations, such as the creation of huge structured-investment vehicles (SIVs). Some transactions are for entities held at their historical basis or book value for a price that would never hold water with any trader. These situations are all examples of pure accounting rigidity.

Another institutional rigidity involving credit ratings relates to certain restrictions placed on institutions that invest in bonds. Sometimes a restriction specifies that a fund may hold only investment-grade bonds. If you're an investment manager holding General Motors Corp. (GM) or Ford bonds, and if those bonds drop below investment grade, you might love holding them or even want to buy more of them. But if you are not allowed to hold bonds rated below investment grade, you don't have any choice: Not only can you not buy any more; you have to start liquidating what you have. That's a rigid rule you are bound by.

Institutional rigidities exist and have to be considered, but hedge funds, because they're not regulated, face far fewer institutional rigidities than other institutions such as banks, insurance companies, pension funds, and mutual funds and thus are the natural counterparties to intermediate the constraints these rigidities create for these other institutions. We normally think of banks as intermediaries that work face-to-face with their customers for their services. Hedge funds can also perform the intermediation function but do so anomalously through the financial markets instead of face-to-face with the entities that are receiving their intermediation services.

I have a test, at least a mental test, for identifying when the hedge fund transactions are providing intermediation services to entities restricted by institutional rigidity: If the hedge fund says to the entity with the institutional rigidity, "I'm going to tell you how we can make money from you in this trade," and even tells the entity how much money it will make, but the entity nevertheless still makes the trade, then it's not that the entity was a less-well-informed counterpart to a transaction. It's simply that the hedge fund was providing a service for the entity by loosening or eliminating a costly rigidity, and it was willing to pay for the hedge fund to do that. I believe that, because of the profits they have been

making, hedge funds have been acting as intermediaries and providing an intermediation service.[12]

III. What the Future Will Look Like for Retirement Plans

We've followed the evolution of pension plans, from the defined benefit,[13] to the bifurcation of alpha and beta, to the reassembling of alpha and beta. Now I'd like to look at that reassembling and consider what opportunities it will give us for the future.

The functionality of employers' providing pension plans for employees is not going to cease. Indeed, our focus-group studies tell us that, in terms of financial services advice, employees place great trust in their employers—more than in any financial services institution. So it looks as if employers should continue to be in the business of providing life-cycle and retirement services to their employees. How are employers handling that part of their business?

By default, plan sponsors are going into DC plans, because they're all getting out of DB plans. But I don't believe that's the end of the story, because I don't believe that the DC plan is a long-term solution as it's typically constituted now. Converting to DC plans does address the concerns of the plans' sponsors; it gets rid of the sponsors' balance-sheet risk and enables them to better contain and calculate their costs as a percentage of wages, which is what they were looking for. But the problem of allocating savings and planning investment portfolios imposes more risk on individuals. It calls on them to make complex and important choices that they didn't have to make in the past, that they're not equipped to make now, and that I doubt that they could be educated to make in the future.

Whether someone is a brain surgeon, a professor, or an auto–assembly-line worker, he or she is now expected to solve a very complex, intertemporal financial optimization problem. What was that DB promise again? *That the retiree would have enough income in retirement to maintain the standard of living he enjoyed in the latter part of his work life.* That's a short explanation that's not a bad approximation of what you might say to twenty-five-, thirty-five-, forty-five-, or fifty-five-year-olds as a likely expectation of what they will want when they do retire. You notice that the statement doesn't mention projected rates of return, and it doesn't address the skill sets needed for understanding risk-return frontiers. The promise is only to solve the problem, and this is what most people will want, a pension solution.

What we're doing for people with DC plans now is giving them a little advice, but basically we're telling them—whether they're brain surgeons, professors, or assembly-line workers—that they are primarily responsible for solving an intertemporal problem like the one in the following box:

The Problem the Employee Has, but Doesn't Want to Deal With

- You're going to be making/receiving contributions or making contributions proportional to your wages for the next twenty-two years.
- You're going to have to invest those monies in something.
- At the end of twenty-two years, you'll have *some* amount of money, but, by the way, it's not the *amount* of money that matters; it's that you want a certain standard of living in retirement, and so you want a *flow* of inflation-protected income that will enable you to maintain that standard of living.

They will have to deal with the problem of not knowing whether their retirement money will be enough or not—even if they know how much money they will have at the end of twenty-two years—because that's going to depend on what the real interest rates are. We've all seen TIPS[14] at 4 percent and then 2 percent; there's a big difference between investing money at 4 percent and at 2 percent. The amount needed to provide income for life will also depend on mortality/longevity tables at the time of their retirement.

So people will have all of those needs and all of those risks, and we're expecting them to solve the problem of optimal accumulation and execute the prescription of that solution. I just don't think that expectation is reasonable as a long-term solution. It's a little bit like going in for surgery, and, as you're being wheeled in to the operating room, having the surgeon look down at you and ask, "Sir [or Madam], would you rather have seven sutures or sixteen?" That seems the sort of thing that the surgeon might be better prepared to answer than the patient. But that's analogous to what's happening with DC plans in the current environment.

People really liked the patterned simplicity of the DB, and they just want to maintain their standard of living. Now we are in danger of repeating the mistake employers made: treating expected returns as risk free over long horizons and undervaluing the cost of DB plans, because we

figured on a 9 percent return when we should have figured on 5 percent. With DC plans, particularly for the longer time horizons, financial advisers too often treat the equity returns very much as if they're darned near risk free. If we don't correct that, we risk having some of the same problems with the DC plans that we had with the DB plans. But this time the consequences of overly optimistic return assumptions and underestimation of the risks to get them will fall directly on individuals and not on corporations, who can at least attempt, even if painfully, to make up the legacy costs.

What do I think is the next-generation solution? I think the next-generation solution is to recognize first that there is a problem that cannot be solved by simply assuming higher returns with the same risk. You can't build an entire retirement system with an essential pillar for success being the creation of alpha. There's a limited amount of alpha out there, and a viable pension solution must be scalable to millions of participants investing trillions of dollars. No matter how talented you are, you can't prudently assume you can continue to create significant alpha—even if you're doing it today—for the next thirty years. That's too aggressive a statement to build a system around.

What are the alternatives? Individuals will have to choose among the following:

- Reducing their consumption during their working years to support the benefits they want postretirement
- Reducing the term of their retirement and working longer
- Lowering the benefits they will receive in retirement

That's it. Those are the choices, and those choices must be part of the process of managing these funds, and there's nothing magical about it. Remember my earlier analogy? When you have Bentleys and you are giving them away at $30,000 apiece, it's great. But once you recognize that the Bentley is a $200,000 vehicle, you're going to have to come up with $200,000 if you still want to buy a Bentley. I don't know what you think, but I think that a lot of people who were quite happy to get the Bentleys at $30,000 will be willing to reconcile themselves to Toyotas, if the price tag on the Bentley is $200,000.

Our job as pension-solution providers and asset managers is to find a mechanism for getting people to be engaged in this part of the process, to reconcile themselves to the realities, and to adjust their positions to arrive at an optimum solution.

We have to improve efficiency in extracting value fully from every asset. Housing, for example, the single biggest asset for most American

baby boomer households, has to be utilized much better than it is now. Essentially, the house is a very special asset at retirement, because it serves two functions: First, it provides prepaid consumption, just like an annuity. If the house is the one the person wants to retire in, ownership is the best hedge for the housing part of his consumption. Using an "average" housing index to calculate your retirement needs won't work if you want to live in Cambridge, or Boston, or New York City. You have to use the housing index for the region you want to live in; that's why ownership of the house is a very good hedge.

In addition to that, however, the house has embedded in it additional value, which, with maintenance, it will have. The trick is to strip out of the house that part of its value that is not prepaid consumption and to utilize that in an efficient way either to add to retirement benefits or, if the goal is to leave money to family or charity, to allow more efficient bequests.

The point is that it's not just housing that has to be added in; *all* of a person's assets have to be accounted for and used efficiently. We are no longer in a position in which we can waste assets; we cannot afford *not* to take into account everything that the individual has, in order to design the best retirement plan. That's a key element.

So, back to my question, What will the future look like for retirement plans? I think that what we will see is that this next generation of plans will be the best of the DB and the DC. From the DB side, we will have the simplicity and the focused target on standard of living. From the DC side, we'll have portability, meaningful customization, minimal credit risk to the beneficiaries, and the absence of residual liability or uncertainty to the corporation. This is a major challenge, but it's also a great opportunity. In the U.S. alone, there are 76 million baby boomers about to retire, and the payout-versus-accumulation solution must be addressed. A daunting challenge indeed. But we also have wonderful, modern financial technology with which to execute a solution, and we have the science of finance as we know it and more than three decades of experience in effectively putting the science of finance into practice.

I believe we have some really exciting and rewarding times ahead for financial services. Providing pension benefits is the challenge, but I think we can meet it. It's an engineering problem, not a basic science problem. Every time you have a major new generational change like this one, there are huge opportunities for creating franchises and adding value. How good is it to be able to come home from work and say, "Hey, we really accomplished something for a lot of people: We're giving them a better, more-efficient retirement"?

Chapter Notes

1. It is perhaps no surprise that corporate pension funds typically allocated 60–70 percent of assets to high-expected-return equities.

2. Some advisers believed that because the time horizon of the pension plan was long, the expected return on the portfolio was risk free, at least over the long run.

3. If that cost is calculated for an 18-year duration, that number doubles to $2 real cost per $1 of perceived cost.

4. By the FASB (the Financial Accounting Standards Board).

5. From the pension fund manager's perspective, the firm's assets, debt, and equity.

6. The direction that we're going in today is to strip out more and more risk, not only from pension fund assets, but also from other parts of the firm's balance sheet, to manage financial risk.

7. A descriptive term, not a technical one.

8. Net asset value.

9. I've assumed away any transaction cost.

10. For technical reasons, even if you use true prices only every other week, the estimated standard deviation comes out about the same as using them every week, because the variance estimate treats the two-week move as if it took place in one week. And so it overstates the true volatility by a factor of two for half the weeks and understates it to be zero for the other half, resulting in an approximately correct estimate for the overall sample period.

11. By demands for short covering.

12. Before you think about regulating hedge funds or any other entities, you need to have an understanding of what functions they're performing. If you regulate hedge funds that are functioning as intermediaries and make them just like mutual funds or banks, they won't be able to perform their intermediary function anymore: Once regulated, they would have the same institutional rigidities as the very entities they're working for.

13. I actually saw it written once as "divine benefit," which wasn't too far from the truth.

14. Treasury inflation-protected securities.