
The Future of Retirement Planning

Robert C. Merton

John and Natty McArthur University Professor

Harvard Business School

Boston

With the move to defined-contribution plans, we, the financial services industry, are asking individuals to make complex financial management decisions that they have not had to make in the past and that, for the most part, they are not adequately prepared to make. In addition, I believe we are presenting these decisions in formats that make them difficult for individuals—even those who are generally well educated—to resolve.

I will begin this presentation with a few remarks about defined-benefit retirement plans, particularly how they went wrong and what we can learn from their flaws. I will then discuss defined-contribution plans, which have become the *de facto* alternative to defined-benefit plans. Unfortunately, traditional defined-contribution plans have a number of features that prevent them from being the long-term answer for employer-sponsored retirement plans. Thus, I will discuss a next-generation solution deriving from defined-contribution plans. Finally, I will discuss financial management technology and the tools available today that can be used to address and help solve the shortcomings of current retirement products.

Defined-Benefit Retirement Plans

Most expert observers agree that corporate defined-benefit (DB) plans are on their way out. The trend in that direction was emphasized in particular when IBM announced in early 2006 that it intended to close its defined-benefit plan to both existing and new employees. IBM is an employee-centric, financially strong company with an overfunded DB plan, and yet the DB plan is being shut down. Some observers say that defined-benefit plans have become too expensive for the corporations to maintain; others say they are too risky. I think the simplest explanation for what happened to defined-benefit plans is that they were mispriced, not three or five years ago but from the outset.

For example, assume that the liabilities in a defined-benefit pension plan have the equivalent duration of 10 years and a risk-free rate of 5 percent. Assume, too, that the plan used a blended expected return on the asset portfolio of 9 percent, not risk adjusted (with assets including risky securities). If liabilities that should have been discounted at 5 percent with a 10-year life span are instead discounted at 9 percent, the result is two-thirds of the present value. Thus, for every \$1.00 of cost a corporation is expecting from a plan, the cost is actually \$1.50.

If a corporation is negotiating with its employees and it offers what it mistakenly believes is \$1.00 of benefits that are really worth \$1.50, then employees are likely to choose the benefits offered over cash, even if they do not know the actual value of the benefits. As an analogy, consider a corporate automobile perk that allows employees to choose either a Toyota Camry or a Bentley. Which will they choose? Will the outcome be random? I do not think so. Even if they have no idea of the actual cost of each, most people are likely to pick the \$300,000 Bentley over a \$30,000 Camry, and just so with generous benefits versus cash compensation.

From the very beginning, providers and sponsors should have recognized that the accounting treatment of these plans was systematically underpricing the cost of benefits. Because of this underpricing, I can say with confidence that we will not go through a cycle that brings us back to defined-benefit plans, at least not to plans with such a pricing structure. Defined-benefit plans have some admirable features, and they may be used again, but we will not return to them with these benefits at this price.

Although defined-benefit plans have been underpriced from the beginning, the reason they are being shut down now rather than 10 years ago is path dependent. During the 1990s, the stock market was up 9 out of 10 years. Therefore, funding for such underpriced plans appeared not to be an issue. But the 2000–02 market crash combined with globally falling interest rates changed that unrealistic outlook, which is why the plans are being reconsidered now.

Defined-Contribution Retirement Plans

The use of defined-contribution plans has become the default strategy following the decline in defined-benefit plans. Although defined-contribution plans solve the problem for the plan sponsor by (1) making costs predictable and (2) taking risk off the balance sheet, they place a tremendous burden of complex decision making on the user.

For example, assume the objective function is that employees hope to maintain the same standard of living in their retirement that they enjoyed in the latter part of their work lives. If that is the goal, then a defined-benefit type of payout is quite attractive. In a defined-contribution scenario, however, a 45-year-old will have contributions coming in for 20 years or more and a 35-year-old for 30 years prior to retirement, and each will need to decide the size of these contributions, as well as the types of investments to make with these funds, in order ultimately to provide the required standard of living at the age of 65.

Finding and executing a dynamic portfolio strategy to achieve such a goal is an extremely complex problem to solve, even for the best financial minds. Yet, through the use of defined-contribution plans, the financial services industry is, in effect, asking employees of all sorts—from brain surgeons, to teachers, to assembly line workers—to solve just such a problem. The situation is not unlike that of being a surgical patient who, while being wheeled into the operating room, has the surgeon

lean down and say, “I can use anywhere from 7 to 17 sutures to close you up. Tell me whatever number you think is best, and that is what I will do.” Not only is that a frightening decision for a patient to be faced with, but it is one that most patients are, at best, poorly qualified to make.

The Next Generation of Retirement Planning. Let me turn to what I think might be a good next generation for defined-contribution plans. If we accept that one of the prospects that most frightens individuals is the possibility of outliving their assets, then the objective function of establishing a standard of living in retirement that approximates the standard of living individuals enjoyed in the latter part of their careers is an appropriate one. Furthermore, if we consider the behavior of most participants in defined-contribution plans, we realize that most people do not enjoy financial planning. After all, most participants do not change their contribution allocations after first establishing them. Therefore, considering individuals’ fear of outliving assets and their disinclination to do financial planning, how should the next generation of plans be designed?

First, if the objective function is an appropriate standard of living in retirement, then the plan should be a system that integrates health care, housing, and inflation-protected annuities for general consumption. Health and housing are substantial factors in the retiree’s standard of living that are not well tracked by the U.S. Consumer Price Index or by any other simple inflation index and should be treated as separate components in providing for an overall standard of living. Furthermore, in order to receive a real annuity at the time of retirement, individuals must expect to pay real prices. Thus, during the accumulation period, real mark-to-market prices should be used. But where do we find such mark-to-market prices? Well, we can approximate them. Insurers, in particular, have the expertise to develop them. What I suggest is that, rather than establishing arbitrary interest rates for the long run, plan developers should use actual market prices derived from actual annuities and mortality experience and mark them to market with respect to real interest rates and not to arbitrary projections. For example, if a plan is based on a 4 percent interest rate and the actual rate turns out to be 2 percent, then the retirees will not have the amount of money they had counted on.

In addition, plans need to be portable. They need to be protected against all credit risks, or at least against the credit risk of the employer. Plans also need a certain degree of robustness, and that robustness must be appropriate to the people who use them. Consider another analogy. If I am designing a Formula 1 race car, I can assume that it will be driven by a trained and experienced Formula 1 driver, so I can build in a high degree of precision because I know the car will not be misused in any way. But if I am designing a car that the rest of us drive every day, I have to be more concerned about robustness than a sophisticated level of precision. When designing a car for the rest of us, I have to assume that the owner will sometimes forget to change the oil or will sometimes bang the tires into the curb. I have to

assume that it will be misused to some degree, so its design must be robust enough to withstand less than optimal behavior and yet still provide the intended outcomes. In applying this analogy to financial plan design, one probably should not assume that users will revise their savings rates in the optimal or recommended fashion.

Qualities of Plan Design: Simplicity and Constancy. What I have in mind is a defined-contribution plan that satisfies the goals of employers while also providing the outcomes of defined-benefit plans, which do such a good job of meeting the needs of retirees. Users should be given choices, but the choices should be ones that are meaningful to them, not the choices that are typically given today, such as what mixture of equities and debt to include in a portfolio. I do not think such choices are helpful for most people.

To use the automobile analogy again, we should be designing plans that let people make their decisions based on a car's miles per gallon, a factor that makes sense to them, rather than an engine's compression ratio, which is a degree of information that most people cannot use effectively. We need to design products that are based on questions that most people find reasonable, such as the following: What standard of living do you desire to have? What standard of living are you willing to accept? What contribution or savings rate are you willing or able to make? Such questions embed the trade-off between consumption during work life and consumption in retirement, and they make more sense to people than questions about asset allocation—or compression ratios.

Besides creating a simple design with only a handful of choices—but choices that are relevant—we need a design that does not change, at least in the way that users interact with it. An unchanging design leads to tools that people will be more likely to learn and use. In fact, a design that is unchanging is almost as important as a design that is simple.

For example, I have been driving for almost 50 years, and during that time the steering wheel in cars has not changed, even though automobile designers could have replaced steering wheels with joysticks. They have been careful to keep the car familiar so that users like me do not have to relearn how to drive each time we buy a new car.

The design of the accelerator is also emblematic of this constancy in design. Depressing and releasing the accelerator requires the same action and provides the same tactile experience that it did 50 years ago. But the technology triggered by the accelerator is entirely different today. Fifty years ago when a driver pressed on the accelerator, that action actually forced metal rods up to the carburetor, opening up passages to allow air and gas to mix and combust and thus send more energy through the engine. Today, the tactile experience is the same for the driver, but the accelerator is not moving metal rods. The processes activated by the accelerator are now electronic. And yet, automobile manufacturers have spent large sums of money making that accelerator feel the same as it did 50 years ago.

The lesson to be learned is that something simple and consistent is easier for people to learn and remember than something complicated and changing. The goal is to be innovative without disturbing the user's experience because planning for retirement is a complicated matter that should not be made more difficult by providing tools that are difficult to use.

Let me return to my automobile analogy. Driving a car is a complex problem. If I wrote down all the information needed to operate a car so that a driver could go from the financial district in Boston to Logan International Airport, I would have a tome full of instructions. It would have to explain the use of the wheel, the gearshift, the accelerator, the brakes, the mirrors, the turn signals, and more. Just getting the car in motion and onto a busy thoroughfare is a complicated coordination problem. Getting to the airport is another level of complexity altogether. And the journey itself is filled with uncertainties. The driver must be alert at all times because, for example, a pedestrian may try to cross the street against the light or a portion of the route may be closed for repairs, and the driver must be prepared to react to each of these uncertainties.

The trip to the airport is difficult enough as it is, but what if the driver is told at the beginning of the drive to the airport, "You must aim the car in the right direction at the start of your trip. After that, you cannot turn the wheel." Knowing the complexities involved in the trip ahead, such constraints make it almost inconceivable that the driver will reach the destination in a satisfactory manner. And yet, most of the models that are used to develop defined-contribution plans implicitly assume that numerous decisions are fixed. That is not an optimal design at all.

We must, therefore, design a system that is user friendly, one that people, given time, can become familiar with and thus willing to use—a system in which the designers do the heavy lifting so that users need only make lifestyle decisions that they understand and that the system then translates into the investment actions needed to achieve the users' goals. The optimal strategies of the system should guide users to arrive at their target retirement goals smoothly. The system will maximize the prospects of achieving a desired standard of living subject to a risk constraint of a minimum life income amount in retirement, but optimization is not simply about ensuring a desired level of retirement income but also about the efficiency or effectiveness in achieving that goal. Just as it is possible to save too little for retirement, it is also possible to save too much and face the regret of forgone consumption opportunities during the many years *before* retirement. Despite these complexities, I am optimistic that such systems are doable, not with futuristic tools but with technology and tools that are available today.

How do I think this next generation of defined-contribution plans will be developed? For one thing, I foresee them developing as corporate plans through plan sponsors because, although the defined-benefit plans are a legacy, I believe

employers will continue to provide retirement assistance in some manner, whether that assistance comes in the form of a 403(b) or a 401(k). One important role employers can play is that of gatekeepers.

Despite the doubts that are sometimes expressed by employees about their employers, when it comes to retirement planning and life-cycle products, people tend to trust their employers far more than they do third-party financial service providers. And employers, despite the criticism sometimes aimed at them, generally want the best for their employees. So, employers can perform a crucial function as reliable gatekeepers when it comes to providing retirement products for their employees.

Technology and Tools for Creating Products

The paradox of the type of system I have just described is that the simpler and easier it is for retirees to use, the more complex it is for its producer. The dynamic trading and risk assessment needed for the next-generation plan require sophisticated models, tools, and trading capability, none of which needs to be explained to the individual.

Interestingly, the mean–variance portfolio model is still the core of most professional investment management models, even for sophisticated institutions. Certainly, it has been updated since its first use in the 1950s, but it is a tribute to Harry Markowitz and William Sharpe that it is still at the core of thinking about risk and return in practice. But to design the next generation of retirement products, designers must consider explicitly some of the other dimensions of risk.

Human Capital. The first dimension is human capital, and the response to include it may seem obvious. But it becomes less obvious how it should be done the more closely it is observed. For example, assume that a university professor and a stockbroker have the same present value of their human capital and the same financial capital. Their risk tolerance is also the same. When deciding which of the two should hold more stocks in their portfolio, most people intuitively respond that the stockbroker should. After all, stockbrokers typically know a lot more about stocks than professors do. But if we consider their situations more closely, we realize that the stockbroker's human capital is far more sensitive to the stock market than the professor's. Therefore, to achieve the same total wealth risk position, the stockbroker should actually put less of his or her financial wealth into stocks. Most models today take into account the value of human capital, but few consider the risk of human capital or how human capital is related to other assets, and that situation needs to change.

Wealth vs. Sustainable Income Flows. The second dimension that needs to be considered is the use of wealth as a measure of economic welfare. To illustrate, consider two alternative environments faced by the individual: One has assets worth \$10 million; the other has assets worth \$5 million. The environment with \$10 million can earn an annual riskless real rate of 1 percent; the one with \$5

million can earn an annual riskless real rate of 10 percent. Which environment is preferable? Of course, if all wealth is to be consumed immediately, the \$10 million alternative is obviously better. At the other extreme, suppose the plan is to consume the same amount in perpetuity. A few simple calculations reveal that the \$5 million portfolio will produce a perpetual annual real income of \$500,000 and that the \$10 million portfolio will produce only \$100,000. So, with that time horizon for consumption, the \$5 million environment is equally obviously preferable. The “crossover” time horizon for preference between the two is at about 10 years. Thus, we see that wealth alone is not sufficient to measure economic welfare.

How many advice engines, even sophisticated ones, take this changing investment-opportunity environment dimension into account? Many such engines quote an annuity (i.e., an income amount) as an end goal, but in doing so, they take an estimated wealth amount and simply apply the annuity formula with a fixed interest rate to it, as if there were no uncertainty about future interest rates. In other words, they do not distinguish between standard of living and wealth as the objective. Sustainable income flow, not the stock of wealth, is the objective that counts for retirement planning.

Imagine a 45-year-old who is thinking in terms of a deferred lifetime annuity that starts at age 65. The safe, risk-free asset in terms of the objective function is an inflation-protected lifetime annuity that starts payouts in 20 years. If interest rates move a little bit, what happens to the value of that deferred real annuity? It changes a lot. If I report the risk-free asset the way typical 401(k) accounts are reported—namely, as current wealth—the variation reported in wealth every month will be tremendous. But if I report it in annuity (or lifetime income) units, it is stable as a rock.

Peru has developed a Chilean-type pension system. A large percentage of the assets—between 40 and 60 percent—are held in one-year (or less) Peruvian debt, with limited international investment. Such a structure does not make much sense for a pension plan. For one, the duration of the bonds should be considerably longer. But every month, the balance is reported on a mark-to-market basis to all plan participants. Imagine the communication challenge of investing in a bond with a 40-year duration, instead, and reporting the resulting enormously volatile monthly balance and explaining why it is actually risk free.

How plans are framed and how their values are reported (wealth versus annuity income units) is thus not trivial. The proper unit of account selected is essential for conveying what is risky and what is not.

Prepackaged Liquidity. Derivative securities can be designed to replicate the payoffs from dynamic trading strategies in a retirement plan. This is done by, in effect, running the Black–Scholes derivation of option pricing “backwards.” Thus, instead of finding a dynamic trading strategy to replicate the payout of a derivative, the financial services firm creates a derivative that replicates the dynamic strategy

desired and then issues that derivative as a prepaid liquidity and execution contract for implementing the strategy. As an example, the dynamic trading strategy for which such prepackaged trading liquidity can be created might be a systematic plan for changing the balance between equity and debt holdings in a prescribed way over time.

Housing Risk. Housing and housing risk are another important dimension, and reverse mortgages are entirely pertinent to this topic. If one is trying to lock in a standard of living for life, owning the house he or she lives in is the perfect hedge. In implementing this aspect of the retirement solution, a reverse mortgage provides an importantly useful tool. A reverse mortgage works within the U.S. tax code to strip out that part of the value of a house not needed for retirement-housing consumption without putting the user at any leveraged risk with respect to the consumption of that house. It is a practical way to decompose a complex asset and use the value to enhance one's standard of living in retirement. It can also be a far more efficient way of creating a bequest than holding onto a house and leaving it to heirs. After all, one does not have to be an expert to know that it is probably far from optimal bequest policy, from the point of view of the heirs' utility, to receive the value of the house as a legacy at some uncertain time in the future—perhaps next year, perhaps in 30 years. I am hopeful that this market will continue to grow rapidly in size and efficiency.

Behavioral Finance and Regret Insurance. For those who believe in its findings, behavioral finance also belongs in the design of life-cycle products. As an example, consider loss aversion, or fear of regret: It appears that loss aversion dysfunctionally affects investors' choices. It inhibits them from doing what is in their best interests. How might we mitigate this problem? Is it possible to create a new financial product, called "regret insurance?" If such a thing is possible, what would it look like?

Consider the following scenario. Assume that a person is broadly invested in the stock market but, for some rational reason, decides to sell. The investor, however, fears that immediately after she sells, the market values will rise. She is frozen by her fear of regret, the regret of selling too low and missing an opportunity to enhance her assets. Fortunately, she can mitigate this situation by purchasing regret insurance. In this case, she buys a policy that guarantees the sale of her stock portfolio at its highest price during the following two years. After two years pass, the investor and the insurer will examine the daily closing price for the portfolio, and the insurer will buy the portfolio for its highest daily closing price during the two years. For the cost of a set premium, the investor is guaranteed an absolute high and is thereby freed of uncertainty and the likelihood of regret.

Such insurance can work for buyers as well as sellers. Suppose an investor wants to buy into the stock market, but he fears that prices will fall after the purchase and he will miss out on better prices. To mitigate his regret, he purchases an insurance policy that allows him to buy the market at the lowest price recorded during the previous two years.

Some people might say that this idea of regret insurance sounds too complicated to produce. How, they might ask, would an insurer determine the risk and then establish a reasonable price? I would submit that such products are already being used in the form of lookback options, which provide exactly the kind of insurance just described. In the exotic options industry, which is quite large, lookback options are frequently issued, which illustrates my general point that the technology and the mathematical tools are already in place to develop the next generation of retirement products. The learning-curve experiences of nearly three decades of trading, creating, pricing, and hedging these types of securities are in place for someone entering the retirement solutions business. It is simply a matter of using market-proven technology in a way that it is not now being used.

Concluding Illustration

One can see from the previous hypothetical example how the identified dysfunctional financial behavior induced by behavioral regret might be offset by the introduction of a well-designed financial product (regret insurance). And if successful, the impact of that cognitive dysfunction on an individual's financial behavior and on equilibrium asset prices can be offset. Note that this change occurs not because of "corrective" education or other means of modifying the individual's internal behavioral makeup but, instead, because an external means is introduced that causes the "net" behavior of the individual to be "as if" such a correction had taken place.

I want to close with a personal, real-world example that illustrates the same dynamics of interplay between the cognitive dissonance of the individual and the corrective effect of the creation and implementation of a financial product or service designed to offset the distortions in financial behavior that would otherwise be obtained, in this case with respect to efficient refinancing of housing mortgages, instead of regret.

In 1999, I took out a mortgage on my apartment, although I do not remember what the interest rate was. Three years later, the same broker who handled my mortgage called me and offered to reduce my mortgage payments by \$400 a month. The offer sounded too good to be true, so I asked what the closing costs would be. He replied that the lender would cover all the closing costs. I then surmised that there must be an embedded option to refinance in my mortgage and that now the lender was trying to get that option out of the mortgage by its generous refinance offer. But the broker assured me that the new mortgage would give me the identical right to

refinance whenever I wanted. Furthermore, the lender was not extending the payment period, and all the other terms of the old mortgage would remain intact, except that I would now be paying \$400 less per month. Even though the deal sounded too good to be true, he convinced me that it was on the level, so I agreed to the refinancing. He came to my office, we signed and he notarized the contract (without my attorney being involved), and the deal has been just as beneficial as he had said.

My guess is that the broker had been given incentives to monitor mortgages like mine for possible refinancing because if he did not get to me, a competitor would. Better to cannibalize your own business by pursuing refinancing than to have the business taken away altogether. Furthermore, my mortgage was probably sold into the capital markets, so his employer, as the originator, would not lose. Certainly, this supposition does not go counter to the way the world works, and thus I ended up being a beneficiary of the competition of the system.

The point of my story is that I turned out to be an excellent illustration of behavioral finance in action. After all, how can someone who does not know the interest rate on his mortgage determine whether he should optimally refinance it? But because of the way the market has developed, the same company that gave me the mortgage gave me a better deal at no cost. I thus ended up behaving like Rational Man in refinancing my mortgage but not because I became “educated” about optimal refinancing models (which I already knew), learned what my interest rate was (which I still do not know), and then optimally exercised. Instead, innovation of financial services together with technology for low transaction costs and market competition allowed me to act “as if” I had. In the process, capital market prices for mortgages were being driven closer to those predicted by the efficient market hypothesis of neoclassical finance. The next generation of retirement products will surely be designed to accommodate and offset such typically suboptimal human behavior.

BIBLIOGRAPHY

Merton, Robert C. 2003. “Thoughts on the Future: Theory and Practice in Investment Management.” *Financial Analysts Journal*, vol. 59, no. 1 (January/February):17–23.

———. 2006. “Allocating Shareholder Capital to Pension Plans.” *Journal of Applied Corporate Finance*, vol. 18, no. 1 (Winter):15–24.

———. 2006. “Observations on Innovation in Pension Fund Management in the Impending Future.” *PREA Quarterly* (Winter):61–67.

Merton, Robert C., and Zvi Bodie. 2005. “The Design of Financial Systems: Towards a Synthesis of Function and Structure.” *Journal of Investment Management*, vol. 3, no. 1 (First Quarter):1–23.

Question and Answer Session

Robert C. Merton

Question: How do you see us moving from defined-benefit plans to something more sophisticated than today's simple defined-contribution plans?

Merton: Most companies want to provide good benefits for their employees. After all, providing efficient benefits is an effective way to pay people, besides its reflecting on a company's reputation. The compensation system of the company is a key strategic issue, and benefits are becoming an ever-larger element of compensation.

That is why such decisions are moving from the CFO to the CEO. They are strategic, and they have considerable implications for the success of the company. Furthermore, companies do not want "smoke and mirrors" for solutions. They do not want to be told that a plan solution will take care of retirement by earning equity market returns virtually risk free, at least over the long run, or by beating the market itself by 1, 2, or 5 percentage points. There are no simple-fix "free-lunch" solutions. That is how we got into trouble with the defined-benefit plans.

Employees, too, need to be told to adjust their expectations. Get used to driving a Camry, not a Bentley, or be prepared to spend \$300,000 to get the Bentley instead of \$30,000 for the Camry. Contribution rates will be whatever is required to achieve goals, and individuals will have to make choices—for example, accept lower wages and higher contribution rates or plan to pay supplemental amounts for additional retirement benefits. They should also expect to make some substitution between the level of retirement consumption and consumption in their working years.

Question: You mentioned that mean-variance is still the norm for modeling, yet the new products you describe have returns that are not normally distributed. Do you believe that many future financing concerns will be addressed by these complex, nonnormally distributed products?

Merton: Yes. But individuals do not need to be aware of this sea change. Individuals do not have to understand the nonlinearities and complexities of option investment strategies or regret insurance. All they are doing is buying a contract that allows them to achieve their targeted replacement of income.

Like the driver of an automobile, individuals do not have to know how the product works. But someone has to know because someone has to be a gatekeeper. Such products cannot be black boxes into which money is poured. But the surgeon, the teacher, and the assembly line worker will not be doing the due diligence. If the plan sponsor is not the one to look under the hood and find out how a product works and decide that it is legitimate, then some other similar mechanism must be established.

Mean–variance portfolio theory is, as I said, at the core of what’s done with asset management for personal finance, asset allocation, mutual funds, and so forth. But the technology to do risk assessment, valuation, and the dynamic strategies I have alluded to, including trying to replicate a 20-year deferred real annuity that is not publicly traded, is market-proven technology that is used every day by Wall Street firms, fixed-income trading desks, capital markets groups, hedge funds, and so forth. Such financial technology is not something that I have plucked off a professor’s idea list. It is in use, and more of it is being created all the time. What I am talking about is a new application of market-proven technologies, and that is why I am confident that the products that I have mentioned are implementable. To be truly an effective solution, the design of such systems, however, must be scalable and cost-effective.

Question: Are you proposing that this new generation of individual retirement investors do not need to understand the risks involved in their plans?

Merton: I am not opposed to people being informed about the investments in their plans, but I think they ought to be informed about things that are useful to them. Disclosing the details of financial technology to nonprofessionals is unlikely to make them any better prepared. What they need to know is that gatekeepers have been established to assure the quality of a plan’s design.

Question: So, you believe a gatekeeper has to be established? In the defined-benefit scenario, corporations did not play that role.

Merton: Plan sponsors did play that role in defined-benefit plans, although not always well. They are still playing that role, and I think they view themselves as such. I do not think they are walking away and closing their defined-benefit plans and telling their employees to go open an IRA.

Corporations have a certain fiduciary duty. Any company that creates a 401(k) plan needs to assure that it is not bogus. If it is, the company is responsible. The company does not guarantee returns, but it must perform due diligence to assure that the plan is sound. One of the strengths of a defined-benefit plan is that it is managed in-house, and the company is responsible for the payouts promised. Defined-contribution plans are outsourced, so plan sponsors have to be more diligent. The one thing I know that does not work is to send surgeons, teachers, and assembly line workers back to school to teach them about duration, delta hedging, and other financial technical details so that they can have a retirement account. Someone else has to take that role. The plan sponsor is probably the natural gatekeeper.

Question: If we need some instruments that are currently not traded, is there a role for government to issue, for example, longevity bonds, so we can derive more information on that type of instrument?

Merton: Absolutely. The creation of TIPS (Treasury Inflation-Protected Securities) and their equivalents was an enormous event, especially for those of us who lived through the big inflation problems with retirement accounts in the late 1970s and early 1980s. As far as I know, only one company—I think it was Aetna working with the Ford Motor Company pension program—wrote an index instrument back then, but it capped the protection at a specific level. Having government-sponsored TIPS and the inflation swap markets that have developed around them means that there will be a lot of ways to develop underlying instruments. Government can play a role.

Certainly, it would be delightful if we had some way of trading longevity efficiently and observing the pricing functions for it. We are not yet there. Still, given the existing environment, I think it is better to use the best estimates from the markets on the cost of longevity risk. In the case of inflation indexing, even though long-dated deferred annuities are not, as far as I know, available at any kind of reasonable prices, one can use well-known dynamic strategies to approximate the returns to such annuities. Such strategies are used all the time in capital market transactions and are done on the other side of these swap transactions to come close to replicating that.

But even if our level of precision for such replications were eight digits, if much of the actual data for other elements of the retirement solution are accurate only to one digit, we do not accomplish a lot by imposing that extra precision. If we get one piece of this problem to a rather precise point, then we know at least one of the elephants in the parlor. For example, in the early phase of retirement accumulation, people do not know what their income will be during the next 30 or so years before they retire. They do not know whether they will be married or divorced. They do not know how many children they will have. They do not know how the tax code will change. We all deal with a host of uncertainties. The idea that we should spend enormous resources making precise one dimension out of all those uncertainties is neither efficient nor cost-effective.

People do, however, need greater precision as they approach retirement because that is when they know much more accurately what they need. They know their income, their marital status, their dependent status. That is also when lifetime annuities become available. Could annuities be more efficient in design and price? No question about it. But they are available, and both design and cost will improve considerably.

Consider another analogy. If you intend to sail from Southampton, United Kingdom, to New York City, you want to aim the boat so that you are not going to Miami, but you do not need a lot of precision at the outset because you can tack as needed to keep a reasonable course. But as you approach New York Harbor, you do not want to be off even by 50 yards. The same concept applies with life-cycle products. A lot of things happen to people during the 30 years leading up to

retirement—many more than they, or we, can predict in advance. What individuals need is a mechanism that allows their retirement planning to adjust to all of these uncertainties as they impinge on them. With the right mechanism, they can adjust their retirement planning and make it more precise as they approach their destination so that, at the end, they can have a lifetime, guaranteed annuity in the amount needed to sustain their lifestyle in retirement.

Now, the guarantee may be, in effect, that of a AA insurance company rather than that of a AAA company. But the guarantee of a AA insurance company is almost always good enough. I am proposing a model that is different from one in which an individual approaches retirement and purchases a variable annuity about which the insurer says, “If the stock market earns 4 percent a year over the Treasury bond yield for the next 15 years, you will be in fat city.” Market risks are a far bigger factor than AA credit risk.

If a stronger guarantee is worth it, institutions may effectively collateralize it. I see such immunization occurring in the United Kingdom, and it may be headed to the United States. Perhaps we will have longevity bonds that are used to immunize against systematic longevity risks. If there is enough worry about the credit of an insurance company, then the annuity need not be run from the general account. The insurer can create SPVs (special-purpose vehicles) in which it can fund the annuities with the right-duration real bonds or with longevity bonds, if those appear.

Although retirement planning today presents a big, challenging problem, it also offers big opportunities. I am excited and optimistic about the prospects ahead. A lot of financial technology exists to help us address the problem.