An Interview with Nobel Laureate Robert C. Merton

Mark P. Kritzman, CFA, and Robert C. Merton

Mark P. Kritzman, CFA, is a founding partner and CEO at Windham Capital Management, Boston. Robert C. Merton is the School of Management Distinguished Professor of Finance at Massachusetts Institute of Technology and the John and Natty McArthur University Professor Emeritus at Harvard University, Cambridge, Massachusetts.

On 7 August 2017, Mark P. Kritzman, CFA, interviewed Robert C. Merton, winner of the 1997 Alfred Nobel Memorial Prize in Economic Sciences, to discuss his student days at Columbia University, Caltech, and MIT; his development of the continuous-time theory of optimal lifetime consumption and portfolio choice as well as his contribution to the development of the option-pricing formula; and his collaboration with Paul Samuelson.

Kritzman: Bob, I'd like to begin by asking you how you came to MIT as a graduate student and, in particular, how you found your way to Paul Samuelson.

Merton: That's a really good question. I was always involved in the stock market all the way through college and then into graduate school at Caltech, where I was getting a PhD in applied math, having gotten an engineering mathematics degree at Columbia. I give this as a predicate to say that I was involved in the markets long before I ever took an economics course.

To fast-forward, I had an awful lot of math before going to Caltech, so I was able to complete my PhD course work in applied math in the first year. I started to think about my thesis. I thought, "What do I find myself doing all the time? I'm always involved in trading and in economics, even though I've had no courses in it." Then, as I thought more and more about it, I said, "This is what I'm interested in." It made no sense, of course, because I was already working on my PhD, but two events happened that I recall very vividly. First, Walter Heller, who had been the head of the Council of Economic Advisers, made a famous claim in the 1960s that macroeconomics had solved the big problems by getting rid of the extremes of depression and big inflation. He said that from then on, it was going to be just a matter of "fine-tuning." I thought that was really amazing. If I went into this field, I could affect millions of people in a good way. That's kind of cool.

Then I went to a Caltech bookstore—again, not having had any economics—and bought a book on mathematical economics, which turned out not to be a good one. I read it and thought to myself, "I might be able to do something in this area." Those two things led me to make a very radical decision, which was to leave Caltech and go into economics.

I applied to nine graduate schools. Eight of them turned me down, including Columbia, where I had gone as an undergraduate in engineering. One took me in. Not only did this school take me in, but it also gave me a full fellowship—and that was MIT. I had a pretty simple decision to make.

Kritzman: Do you have any insight about why MIT chose to accept you?

Merton: Yes, I do. I found this out later. There was a man named Harold Freeman, who was a professor of statistics in the economics department and who had been instrumental in bringing Paul Samuelson to MIT. This was back in the days when the economics department was nothing more than an adjunct to the engineering school. Harold was on the admissions committee. When I had applied earlier for my PhD in applied math, I had applied to two schools: Caltech and MIT. He saw the references written for my mathematics application and recognized the people who wrote them. Apparently, he convinced the committee to take a flier with this guy who had no economics background.

So, I came to MIT and showed up the first day for registration. I was very contrite. I went to the graduate officer who advised me, and I said, "I've read the brochure. I'm going to take micro and macro and economic history." That guy happened to be Harold Freeman. Harold looked at me and said, "If you take those courses, you're going to be so bored, you'll leave here at the end of the term. Go down and take Paul Samuelson's mathematical economics course." I said, "But Professor, that's a second-year course and I've had no economics." He said, "Just go take it." Being a dutiful student, I signed up. That's how I was introduced to Paul Samuelson.

Kritzman: What was he like?

Merton: We could talk all afternoon just about him. I can tell you that I took the course, and one day after class, he asked me if I would be willing to check the math in a paper he had written. I took it home and I stayed up all night working on it. I went through every page more than once and got it all done by the next day. I brought it to his office and tried to be nonchalant. I said, "I have some comments." When I went to class the next day, he asked me, "How would you like to work for me?" That's how I got involved with Paul Samuelson.

When Paul and I got to talking with each other, we found out we had a common interest in the stock market—in particular, warrants. We hit it off. At some point, he said, "I wrote this paper on warrant pricing.[1] It was published back in 1965. I really want to do an equilibrium model, a model where you could derive the pricing." That was our first joint research. Of course, I had not had any of this stuff, but it didn't take me long to learn it.

Kritzman: When you graduated, you stayed at MIT. Isn't that unusual? Most people go on to other schools.

Merton: I had a very productive graduate experience at MIT. I had five chapters in my thesis, and three of them were published before the entire thesis was written. The fourth chapter was published within a year after I graduated. I had a lot of other research ideas going. I was at MIT for three years. I could have finished earlier—but why? I was doing good work.

It came time for the job market. MIT Economics had an absolute rule: You have to leave; it cannot hire its own. So, when it came time to go on the job market, I didn't even think about MIT. And I interviewed only with economics departments, because I didn't know anything about business schools. It never even occurred to me to interview with business schools. I got offers from the University of Chicago, Yale University, Princeton University, and a few other places. As I was thinking about where I should go, Franco Modigliani came to me. He was interested in my thesis work because he viewed my work on intertemporal lifetime consumption as very complementary to his life-cycle hypothesis. He knew what I had done. I had taken his classes. He came to me and said, "How would you like a job at the MIT Sloan School of Management in the finance group?" I said to him, "Well, Professor Modigliani, for one thing, I don't know anything about the business school. I've never taken a finance course, so it would seem a bit of a challenge to become a professor teaching finance." He said, "With your thesis, you'll have no problem."

As I was thinking about where I should go, Franco Modigliani came to me. He was interested in my thesis work because he viewed my work on intertemporal lifetime consumption as very complementary to his life-cycle hypothesis.

I mulled it over. I was living in the tower right next door to the parking lot here at MIT and was all settled in. I had been very productive here. I wouldn't have to move. I could just finish my thesis, go right to work, and not have to do anything. I wouldn't have to interrupt my research. I wouldn't have to move family. I wouldn't have to do anything, so I took the job.

It didn't violate the rules because it was a different department. I had graduated Course 14, and this was Course 15. The powers that be seemed to think that was just fine.

Kritzman: To what extent did you continue your interactions in the economics department?

Merton: Of course, I continued my interactions with Paul Samuelson. We did several papers right into the mid-1970s. The papers we did were good papers. They were published, but they certainly were not his or my own most significant contributions.

I had this idea that I wanted to do intertemporal optimization under uncertainty and also do the lifetime consumption portfolio problem. I did it using a continuous-time model. I showed it to Paul [Samuelson], and he really liked it.

I had this idea that I wanted to do intertemporal optimization under uncertainty and also do the lifetime consumption portfolio problem. I did it using a continuous-time model. I showed it to Paul, and he really liked it. He decided to do a companion paper on discrete time. It was great for me because my paper was published very quickly. I had written the paper in the late summer and early fall of 1968, and it was published in August 1969 in the *Review of Economics and Statistics*. My paper[2] and Paul's paper[3] appeared next to each other as companion papers.

Kritzman: You're best known for the development of the continuous-time theory of optimal lifetime consumption and portfolio choice and for your contribution to the development of the option-pricing formula. How did you become interested in continuous-time finance? Are you the person who brought Itô's lemma to finance?

Merton: Yes. I was always interested in dynamics, optimization, and uncertainty. As I mentioned earlier, I wanted to put them all together, but I didn't know how to do it because I had never seen stochastic dynamic programming. Instead of learning the mathematics and then looking for a problem to use it on, I had the problem and I found the mathematics that I needed. Happily, I had had quite a bit of training in mathematics. I had mathematical maturity, which means that you can teach yourself new math without having to go back to class. I found stochastic dynamic programming, and I thought, "Hey, this works." That's how I did my first paper on dynamic optimization under uncertainty [Merton 1969].

As I studied stochastic processes in my search, I kept running across footnotes about this fellow, Itô. I explored it because I was looking for a way to describe not the expectations of where prices, wealth, and consumption are going but, rather, the actual sample paths themselves—the realizations. I could do the basic portfolio problem without this because stochastic dynamic programming is about expected utility. So, ultimately, you're taking expectations of a function of the outcomes. I really wanted to understand the dynamics of the paths themselves. I started looking around, and I found the Itô calculus. To put it in context, there was another stochastic calculus by a Russian named Ruslan Stratonovich. His calculus, unlike the Itô calculus, satisfied the rules that you and I learned in college. Itô's calculus had a different set of rules. To a physicist, either calculus was fine. But it turned out that for modeling what I was doing, the Itô calculus was the right one to use, and the Stratonovich calculus, which would give you different answers, was the wrong one to use. What Itô allowed me to do was write down the dynamics of any function of a stochastic process of this kind, not just the expectations. This second paper on the optimum consumption and portfolio problem [Merton 1971] was considerably more extensive than my first, and it also became a chapter in my thesis [Merton 1970].

My dad taught me enough about priorities. As far as I know, I was the only one who even thought about using the Itô calculus for describing economic dynamic sample paths. It's the sort of specialized topic you weren't taught even if you studied math extensively at that time.

Doing the modeling in continuous time was a natural for me because I liked the calculus of variation. With the discrete-time approach, you always faced the question, what's the interval between actions? I thought, "There probably aren't any general theorems to be derived for an arbitrary time interval. What about the idea of using the shortest feasible trading interval as the specified interval and modeling it by continuous trading?" Obviously, you couldn't literally do it, but I could justify it as a reasonable approximation to reality because of the nature of the mathematics. I could show that even if you traded once a day, it's very close to continuous in the mathematical sense. What impelled me to do the continuous-time finance was a combination of an intellectual challenge—and it more closely matched the reality of market trading—and the knowledge that discrete-time mathematics has always been difficult to work with.

Kritzman: How did your work in continuous-time finance advance the science of finance as it relates to portfolio theory, the capital asset pricing model, and derivative security pricing?

Merton: The first paper I did was published in 1969. "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case" didn't use the Itô calculus; it didn't need to. It just used dynamic programming.

Lognormal was the prototypical distribution for stock prices at that time. I said, "What if stock prices were all joint lognormal?" When I solved that problem, what came out was that the demand functions for the optimal portfolio were exactly the Markowitz–Tobin demand functions—the mean–variance ones—and it had all the same features. What got me so excited was that I took a more realistic model—multi-period—and didn't have to assume normal distributions, with their prospects for negative prices. The Markowitz–Tobin rules worked for any concave utility function, so I didn't have to assume any kind of particular utility function other than risk aversion. And boom: Out come the Markowitz–Tobin rules, exactly.

I was able to reconcile the practical, intuitive, closed-form-solution optimal portfolio rules of the Markowitz–Tobin mean–variance theory with expected utility theory and a prototypical lognormal return distribution, which doesn't have the negative-price perversities that a normal one has.

I was able to reconcile the practical, intuitive, closed-form-solution optimal portfolio rules of the Markowitz–Tobin mean–variance theory with expected utility theory and a prototypical lognormal return distribution, which doesn't have the negative-price perversities that a normal one has. When you worked it through, it turned out that when you keep constant proportions—as you would—nothing changed in terms of the risky assets' distributions. When you trade continuously and combine lognormals, they aggregate to a lognormally distributed portfolio. Therefore, the Markowitz–Tobin two-fund separation theorem obtains as well.

When I learned about the CAPM, I took my model and put it into an equilibrium context. Of course, I knew what I was going to get: the CAPM. Thus, the CAPM holds over the trading interval but only if you measure returns instantaneously. However, it was certainly not true if, instead, you measured returns over a month or six months or a year or five years. Returns can be linearly related over only one specified time interval, and I chose the minimum interval to be continuous.

When I expanded the model to include stochastically changing expected returns and covariances, I was able to show that the same linear structure of equilibrium expected returns would obtain, but with multiple dimensions of systematic risks. Instead of the security market line of the CAPM, one has the security market hyperplane. In 1975, I published an early overview of my perspective on the contribution of the continuous-time model to finance [Merton 1975].

In sum, I felt that it was cool that by going to a more realistic model, one got simpler, more robust results that worked, which is not what one expects to happen.

Kritzman: Let's turn to your work in warrant and option pricing. Were you doing this before you met Fischer Black and Myron Scholes or knew what they were doing?

Merton: Absolutely.

Kritzman: My understanding is that Fischer and Myron approached the problem from an equilibrium perspective, whereas you thought about it from the perspective of eliminating arbitrage through continuous trading.

Merton: That's approximately true, but let me tell you the story. When I was still a graduate student, I had to come up with an ad hoc model called the "equal return for equal risk model." It was completely ad hoc, but if you solve that model, you get the Black–Scholes formula—but not for the right reasons.

I wrote a paper [Merton 1968] on general no-arbitrage conditions without distributions or theorems about puts and calls. I did a lot of work on that. I was fascinated. I had worked with Paul Samuelson and published a paper [Samuelson and Merton 1969] on an equilibrium model of warrant pricing. I then did convertible bonds the same way, because once you've done the warrants, you realize you can do other contingent securities the same way.

I was working on that when I met Myron. Fischer was at Associates in Finance, with one employee: him. Myron and Fischer had been working on option pricing. Option pricing was a considerably greater challenge than forward or futures pricing because the option has a nonlinear payoff function. There is no static position you can take in stock and cash the way you can in a forward contract that will completely hedge it or replicate it. It isn't doable. There is no static position that can do it.

The brilliant idea they had was, "What if we did a dynamic trading model where we took a position in the stock and cash, a risk-free asset, and the option, and we set the weights so that the portfolio had a beta of zero? We've got three assets, two of them risky. We would adjust the mix to make the overall portfolio zero beta. We would let the thing play out for a period, and then we would change the portfolio to reset it so it has a beta of zero again for the next period." That was the core idea they had. They then wrote that process in discrete approximations. They used a normal distribution approximation over the trading interval.

Kritzman: For convenience.

Merton: For convenience, because over a small interval, lognormal was like a normal. If they could get that equation, they knew from the CAPM in equilibrium that the expected return on a zero-beta portfolio is always equal to the risk-free rate. When you put that equilibrium condition in, that can close the model to derive a partial differential equation for the price. They got the equation by applying the CAPM but didn't know how to solve it.

There's a great irony here. The instant I saw that equation, I knew what it was because if you had ever read Paul Samuelson's paper "Rational Theory of Warrant Pricing" [Samuelson 1965], it is Samuelson's equation exactly, but he simply posited parameters for the expected return on the warrant and on the stock. If you look at the equation and substitute the risk-free rate for those two numbers, that's it!

Kritzman: Paul didn't use the risk-free rate in his formula.

Merton: No, not at all. I'm talking about this to explain the same pricing equation structure. It's interesting because Fischer had said for a long time that he couldn't find a solution to the equation. Had he read Paul's paper "Rational Theory of Warrant Pricing"—which was published years before, in 1965, in the Sloan School journal *Industrial Management Review*—he would have seen both the equation and the answer. Of course, for me, having lived with that paper, the instant I saw that equation, I knew it. Not only did I know how to solve it, but I also knew how to interpret it.

Kritzman: How did you find out about this?

Merton: Myron and I got to know each other. I had met him earlier when I was still a student, because he did the job interviews. And since I was in the same building, we had talked, but I didn't join the faculty until the summer.

At some point along the line, he came to me and said, "Let me tell you what Fischer and I have been doing." They were working by themselves, and I was working by myself. We never saw each other's work. When Myron told me what they were doing, I said to him, "That can't work because the option price is a nonlinear function and portfolio combinations are linear." He then showed it to me. I thought, "How would I think about this?" Of course, I just took my continuous-time model, because they were doing dynamic trading over a short, but discrete, interval. I had the mechanics of the Itô calculus from my portfolio theory research, and so I just took what they were doing and said, "How would you do this if you literally traded continuously?"

Guess what I discovered? Not only did you get rid of the systematic risk, but you also got rid of *all* the risk. It wasn't that you made it zero *beta*. You made it zero *sigma*.

The big difference is that a necessary condition for every equilibrium model is no arbitrage. Any equilibrium model that satisfies the conditions assumed for Black–Scholes has to give you that price because it comes from no arbitrage. It has nothing to do with the CAPM if you assume that you could trade continuously. That's the trade-off. If you literally could trade continuously, that's my model. All the risk goes away!

I went back to Myron and said, "I was wrong. My intuition was wrong. You were right, but for a different reason." Then I pointed out that since it's a necessary condition, I believe it is the more general way to look at it because now your model works for all equilibrium models. I said, "You don't need any of this. You don't need any assumption about which equilibrium model applies. You really don't need equilibrium." That's the contribution I made to derivatives pricing that ultimately earned me my share of the Nobel Prize.

I teach Black–Scholes as a production theory for options or for any derivative, because it doesn't matter whether the payoff is a call or a put or a squiggle. It doesn't change any of the theory. It's just a different boundary condition. The Black–Scholes price is the production cost to the lowest-cost producer, who in my world can trade continuously at no cost. That's the extreme of the lowest-cost producer. The power of the replication says that even if there are no derivatives or options, I can manufacture them because I have a trading strategy that exactly replicates their payoffs. Removing any risk from the hedge portfolio means that it must be exactly replicating the payoffs. The formula for the deltas or for hedging is the production function. It says, "As you go through the process, change your portfolio according to this rule depending on what happens, and out the other end comes a payoff function that is identical to that derivative."

The Black–Scholes price is the production cost to the lowest-cost producer, who in my world can trade continuously at no cost.

Kritzman: Just to try to summarize a bit: Fischer and Myron came up with this formula. They didn't quite know how to solve it. As a convenience, given the very short time interval, they assumed normality, which enabled them to come up with that formula. Is that right?

Merton: Yes. I believe what helped them without the math is that they did the Taylor series expansion of the function, just as we all learned. When they did their expansion, they got a normal variate multiplying a squared normal variate. It's always true that the expected value of that product is zero because of symmetry. They got rid of it when they did the beta part. It dropped out for that reason. They never had to explain the complicated part of Itô's calculus. Where did it go? I wrote a whole paper [Merton 1982] about that, but that's how they got the right answer. It came from the fact that the square of a normal and a normal always have zero expectation.

Kritzman: If they had been more precise in assuming lognormal distributions, they wouldn't have gotten there.

Merton: Yes. I think that they were doing what was very normal practice. No pun intended.

The reality was that we did our papers separately. Their paper[4] was at the *Journal of Political Economy* for a long time, and they kept getting comments like, "This is too specialized for the *JPE*." I believe Merton Miller hammered the *JPE* until the journal accepted it.

My paper[5] was published in the *Bell Journal of Economics and Management Science*. My new colleague, Paul MacAvoy, was the first editor of the journal. He came to me in 1971. He was looking for papers. He said, "If you contribute your paper, you can have as much room as you like and I'll pay you $500." Now, my MIT salary was $11,000, so that helps to scale what $500 was worth. It was not a small amount, especially for a junior faculty member. I said to him, "Under one condition: I get to pick when it comes out. I will tell you in advance when to publish it." And he was fine with that.

Obviously, the big concern that I had—especially with my sensitivity to issues of priority, from my father being a sociologist of science—was not to publish my paper before Fischer and Myron published theirs. People might infer that their work was done after mine because of the order of the publication of the two papers. That's why I delayed publication, and so, when their paper came out in the spring of 1973, mine did too. That was not a coincidence. I felt that that was the best way to do it, and I think it worked well.

Kritzman: Still, that was very generous of you.

Merton: Not generous—it was just the right thing to do.

Kritzman: You have pointed out that the introduction of derivatives poses a challenge to the stability of the CAPM as well as to Steve Ross's arbitrage pricing theory (APT) and that continuous-time trading resolves this stability issue.

Merton: Yes. Let's consider APT, because it's very important work and it's a general model. Suppose you're in an equilibrium and everything fits the model. First, you must specify over what time its linear structure holds. You cannot have it hold over all periods because returns compound. You must pick your period. Second, if anything disturbs that linear structure, the equilibrium will not hold.

Let's say you have a 10-factor model that correctly describes the equilibrium asset prices. Someone issues a warrant or an option. All of them are nonlinear functions. That warrant or option will not fit those 10 linear factors, because although the factors are there, it's a nonlinear mixing of them. It's all jumbled up. If you had a linear structure that worked before, it will no longer work. Hence, the equilibrium model is not robust to such changes.

When you can trade continuously, you can replicate all derivatives' returns by trading the underlying securities. Therefore, the continuous-time model is "closed" under derivatives, which means that you preserve the same equilibrium prices with or without derivatives in the system. The market is already complete with respect to derivatives. Because you could have replicated all those derivatives beforehand, adding the derivatives doesn't expand the space of available payoffs. That's what I mean by saying it's closed. You can create any contract structures you want. You can have new kinds of securities no one has ever seen before—like squiggles—and it won't affect the equilibrium. I don't think any other models satisfy that except the Arrow state preference model, because there's no residual. There's a finite number of states that can happen, and you have that number of securities to cover all states.

When you can trade continuously, you can replicate all derivatives' returns by trading the underlying securities. Therefore, the continuous-time model is "closed" under derivatives, which means that you preserve the same equilibrium prices with or without derivatives in the system.

Kritzman: You brought up Ken Arrow's state preference theory, which assumes many different security markets but very little trading, whereas the reality is lots of trading but few security markets.

Merton: This takes us back to 1953, when Arrow published his "Role of Securities" paper, which I consider one of the fundamental contributions that mark the beginning of finance science, along with Harry Markowitz's work.

Everyone looks at that and says it's a beautiful, elegant theory and that it gives us an idea of what security markets can do and what they can't do. But it's totally unrealistic, because it predicts you're going to have an uncountable number of different markets, a market for every state, a security for every state and time, and no trading. What we observe is relatively few securities—stocks and bonds and so on—and trading all the time.

The replication analysis that I contributed reconciles that. It says that we can have either (1) a market with all kinds of derivatives markets, a derivative for every state that can happen, and thus dozens and dozens and dozens of markets or (2) the trading in the assets underlying all these contracts, with trading in them continuously. We can then replicate all those contracts or manufacture them, so we can make only the ones we need. We don't actually need to have that market in place for all payoffs, because we can manufacture the payoffs we need, as in just-in-time production.

What that showed was that the institutional framework that Arrow used was not descriptive of the institutional world we observe. He chose an institutional setting—understandably, because it's very clean—that is not at all descriptive of the real world.

I showed that there is this alternative real world where you're trading all the time—not because people's risk preferences are changing or because people get new information, but because markets change and we must adjust the portfolio holdings to accommodate the desired payoffs of the strategy. If we're doing this replication, we have to trade, and we know that the replication principle does not depend on expectations or whether we think the underlying assets are under- or overpriced. It's purely manufacturing the payoffs to contracts on those assets. Trading occurs because people are producing these payoffs. It's more efficient to do it this way than to have uncountable markets and no trading.

In summary, the continuous-trading assumption reconciles Markowitz–Tobin and the CAPM with expected utility theory. It resolves the issue of the instability of linear asset pricing models. And finally, it allows us to create Arrow securities.

Kritzman: Bob, let's switch gears. Could you talk a little bit about your father? I think most people know that he was a very famous sociologist, having coined the terms *role model* and *self-fulfilling prophecy*.

Merton: Yes, he loved the title "The Self-Fulfilling Prophecy." Because with it, you didn't have to read the paper. All you needed was to read the title. You knew what it was going to say. "Role Model" was not quite so obvious but of the same genre.

Kritzman: He was the most or, if not the most, one of the most prominent sociologists of his day, so you grew up in this household with a father who was a luminary in academia. What was that like, and what effect did it have on your choice to pursue an academic career?

Merton: It's a fine question. First, I'll mention that we were really best friends for 40 years, right up to the end of his life. I knew him longer, but when I was a younger man, I had to find my own way.

I'll start with things he *didn't* do. One thing is that he never really asked to look at our grades—my grades or my sister's grades. He never pressed us in that way. He didn't have to. We just watched what he was doing.

Kritzman: You understood who he was at the time?

Merton: Yes, my siblings and I understood as kids, and we understood it maybe more a little later, but we understood it pretty early. I was exposed to academics of all kinds—Nobel Prize winners—from an early age just because that's what I grew up with, but the one thing I knew I was not going to be no matter what was a professor. I thought I would probably be an auto engineer. I used to build cars and race them. I loved them. I even worked for Ford in advanced vehicle design a couple of summers when I was in engineering school. I just knew I wasn't going to become a professor, and certainly not in the social sciences. But here I am—half a century later—a professor in the social sciences.

I think the way he most affected me was through his respect for mathematics.

Kritzman: Was he a mathematician himself?

Merton: No, not at all. He was a magician, a professional magician, a very smart guy, analytical—absolutely—but not mathematical. But he had great respect for math, and that was my thing.

I never thought I would end up in a field that was close to his, but if people read my work, they will notice a couple of things. First, they will see that I have written a lot over the last 25 years on applying a functional perspective to understand the endogenous dynamics of institutional change, driven by more efficient institutional designs to perform the functions. I have used his terms *manifest function* and *latent function* over and over again in describing markets. Markets have a manifest function of trading to exchange, but markets also have a very important latent function of providing information from the prices generated by these exchanges in the markets.

Markets have a manifest function of trading to exchange, but markets also have a very important latent function of providing information from the prices generated by these exchanges in the markets.

The other thing was that I never felt intimidated as a student. Impressed? Yes. I looked with awe at what some people could do, but I was never intimidated by even very famous, Einstein-type minds. I had enormous respect for them, but I was never intimidated by them. I think part of that was just familiarity and growing up in an environment where I saw that all of us put our clothes on the same way, and so I always felt comfortable. I don't know why that was, but I have to assume that that was not independent of my experience growing up with my father and the people around him.

Kritzman: Bob, is there anything else you would like to share?

Merton: Yes. Over my career, I've been surrounded at MIT by an incredible set of colleagues: Stewart Myers and Myron Scholes, and then Fischer Black and John Cox, who came later, and of course, Paul Samuelson and Franco Modigliani—to name just a few from the early days. If one has a big head, it gets deflated pretty quickly if you walk around MIT.

There's no substitute for having both very nice people and very talented people around you. What I've said to you describes the way I looked at finance—my vision, not "this is the way the world is or was." So, together with others that you may have interviewed or will, people can get a more complete picture in some sense.

That's the spirit in which I've answered your questions. Obviously, I feel—as you can understand—a little self-centered spending two hours talking about myself. But in talking about reconciling this and doing that—some pretty important things, actually—I hope I have given you what you asked for.

Kritzman: You have. And I have no doubt that the readers of the *Financial Analysts Journal* will have a much deeper appreciation of the evolution of modern finance.

Thank you!

Notes

1. Paul Samuelson, "Rational Theory of Warrant Pricing," *Industrial Management Review*, vol. 6, no. 2 (Spring 1965): 13–39.
2. Robert C. Merton, "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics*, vol. 51, no. 3 (August 1969): 247–257.
3. Paul Samuelson, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, vol. 51, no. 3 (August 1969): 239–246.
4. Fischer Black and Myron Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, vol. 81, no. 3 (May–June 1973): 637–654.
5. Robert C. Merton, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, vol. 4, no. 1 (Spring 1973): 141–183.

References

Arrow, Kenneth J. 1953. "Le rôle des valeurs boursières pour la repetition la meilleure des risques." *Econometrie*, vol. 11. Translated as "The Role of Securities in the Optimal Allocation of Risk-Bearing." *Review of Economic Studies*, vol. 31, no. 2 (April 1964): 91–96.

Black, Fischer, and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, vol. 81, no. 3 (May–June): 637–654.

Merton, Robert C. 1968. "Restrictions on Rational Option Pricing: A Set of Arbitrage Conditions." Mimeo, Massachusetts Institute of Technology (August).

——. 1969. "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case." *Review of Economics and Statistics*, vol. 51, no. 3 (August): 247–257; reprinted as Chapter 4 in Merton, *Continuous-Time Finance* (Blackwell, 1992).

——. 1970. "Analytical Optimal Control Theory as Applied to Stochastic and Nonstochastic Economics." PhD dissertation, Department of Economics, Massachusetts Institute of Technology.

——. 1971. "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." *Journal of Economic Theory*, vol. 3, no. 4 (December): 373–413; reprinted as Chapter 5 in Merton, *Continuous-Time Finance* (Blackwell, 1992).

——. 1973. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science*, vol. 4, no. 1 (Spring): 141–183.

——. 1975. "The Theory of Finance from the Perspective of Continuous Time." *Journal of Financial and Quantitative Analysis*, vol. 10, no. 4 (November): 659–674.

——. 1982. "On the Mathematics and Economic Assumptions of Continuous-Time Financial Models." In *Financial Economics: Essays in Honor of Paul Cootner*, edited by William F. Sharpe and Cathryn M. Cootner. Englewood Cliffs, NJ: Prentice Hall; reprinted as Chapter 3 in Merton, *Continuous-Time Finance* (Blackwell, 1992).

——. 1992. *Continuous-Time Finance*, rev. ed. Oxford, UK: Blackwell.

Samuelson, Paul A. 1965. "Rational Theory of Warrant Pricing." *Industrial Management Review*, vol. 6, no. 2 (Spring): 13–39.

——. 1969. "Lifetime Portfolio Selection by Dynamic Stochastic Programming." *Review of Economics and Statistics*, vol. 51, no. 3 (August): 239–246.

Samuelson, Paul A., and Robert C. Merton. 1969. "A Complete Model of Warrant Pricing That Maximizes Utility." *Industrial Management Review*, vol. 10, no. 2 (Winter): 17–46; reprinted as Chapter 7 in Merton, *Continuous-Time Finance* (Blackwell, 1992).