SESSION 1: BODIE PRESENTATION

The Theory of Optimal Life-Cycle Saving and Investing

Zvi Bodie
Norman and Adele Barron Professor of Management
Boston University
Boston

Jonathan Treussard
PhD Candidate
Boston University
Boston

Paul Willen
Senior Economist, Research Department
U.S. Federal Reserve Bank of Boston
Boston

Life-cycle saving and investing are today a matter of intense concern to millions of people around the world. The most basic questions people face are

- How much of their income should they save for the future?
- What risks should they insure against?
- How should they invest what they save?

Many of these questions are answered for people—in whole or in part—by government, trade unions, employers, and other institutions. In this paper, we argue that economic theory offers important insights and guidelines to policymakers in government, to the financial services firms that produce life-cycle financial products, to the advisers who recommend to their clients which products to buy, and to educators who are trying to help the public make informed choices.

The literature is vast and complex, and we will not attempt to survey or summarize it all in this paper. Instead, we lay out the basic analytical framework using a few relatively uncomplicated models and focus on several key concepts. This analytical framework could serve as a valuable guide to financial services firms in helping them to develop and explain products in terms that are understandable to the layman.

In the first section of this paper, we introduce the life-cycle model of consumption choice and portfolio selection. We emphasize the central role of consumption in life-cycle planning, and we highlight the use of financial assets as means to

©2007, 2008, The Research Foundation of CFA Institute

transfer consumption from points in the individual's life cycle when consumption has relatively little value to points when consumption has relatively more value. The second section highlights five concepts from the life-cycle model that are directly relevant to the practice of life-cycle planning: (1) the notion of a lifetime budget constraint, (2) the relevance of contingent claims in life-cycle planning, (3) the trade-off imposed by varying costs of consumption over one's lifetime, (4) the role of risky assets, and (5) the asset allocation decision over the life cycle. In the third section, we recognize the complexity of life-cycle planning and the relevance of financial frictions at the individual level, and accordingly, we argue that it is the role of specialized firms to engineer and deliver life-cycle products that meet the needs of households. The final section contains our concluding remarks.

Theoretical Introduction

The starting point for analysis of life-cycle portfolio choice is a model of the evolution of an investor, which one can think of as an event tree.[1] **Figure 1** illustrates an event tree for a fictional investor who lives for three periods: youth, prime earning years, and retirement. In addition to aging, events occur that affect the investor. In Period 2, he earns either high or low income, and in Period 3, he enjoys either good or bad health. Figure 1 shows income and, in retirement, expenses associated with the different outcomes. We assume that our investor earns no income in retirement and faces no health expenses before retirement. A financial plan, in this context, tells an investor how much to save or borrow, how to invest any savings—not just today (in Period 1) but in the future (in Period 2)—and how much the investor should withdraw in retirement (Period 3). That plan could also depend on contingencies: Along the high-income path, our investor may want to save more; along a bad health path, our investor may want to withdraw less to prepare for high bills.

Suppose, for simplicity, that the only investment opportunity is to save or borrow at 0 percent interest. The lines labeled "advice" in Figure 1 reflect a simple proposed financial plan: Save $10,000 a year when young, save $20,000 a year during prime earning years, and withdraw $50,000 a year in retirement. It is easy to verify that this plan works. Savings at retirement equal $750,000, which, divided up over the remaining 15 years, allows withdrawals of $50,000 a year.

How good is this proposal? Standard approaches to financial planning would focus on whether the investor could afford to save enough or whether the $50,000 would be enough in retirement. What does the life-cycle theory say about this proposal? We distill the following three principles from the life-cycle approach.

[1]In finance, the event tree has become a workhorse tool—most importantly, the Cox–Ross–Rubinstein (1979) binomial model.

Figure 1. The Life-Cycle Model: An Event Tree

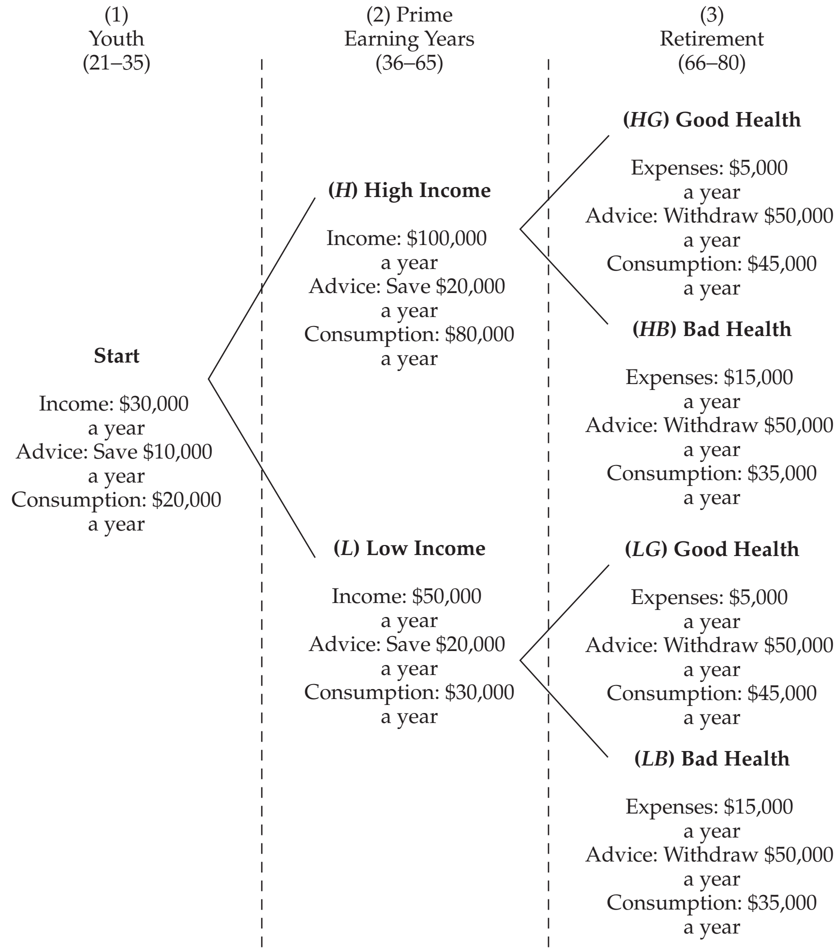

Principle 1: Focus not on the financial plan itself but on the consumption profile that it implies. In this example, we can easily calculate consumption (shown in Figure 1) because it equals income less savings during working years and withdrawals less health expenses in retirement.[2]

Principle 2: Financial assets are vehicles for moving consumption from one location in the life cycle to another. Suppose, for example, our investor wanted to increase consumption in youth. By reducing savings in youth and leaving it unchanged in prime earning years, our investor can transfer consumption from retirement to youth.

[2]This insight goes back to Fisher (1930) and Modigliani and Brumberg (1954, 1979).

©2007, 2008, The Research Foundation of CFA Institute

By reducing savings in youth and raising savings in middle age, our investor can transfer consumption from prime earning years to youth.[3]

Principle 3: A dollar is more valuable to an investor in situations where consumption is low than in situations where consumption is high. In Figure 1, for example, the life-cycle model says that if we offered to give a dollar to our investor but said he must choose when he wanted it, our investor would want the money in youth, when his consumption is lowest. The law of diminishing returns is at work here: An additional dollar is a lot more valuable to a recent college graduate than to a 55-year-old executive.

So, what does the life-cycle model tell us about the advice in Figure 1? Looking at the implied consumption over the life cycle, we notice huge variations. According to our third principle above, a dollar is much more valuable when consumption is low than when it is high. And so we can improve on this plan by trying to move consumption from situations with high consumption to situations with low consumption.

- Consumption in youth ($20,000) is much lower (and thus more valuable) than it is on average in prime earning years ($55,000). By saving less in youth and more in prime earning years, our investor could transfer consumption from low-value situations to high-value situations and make himself better off.
- Consumption in situations *HG* and *HB* ($40,000 on average) is much lower (and thus more valuable) than in situation *H* ($80,000). By saving more in situation *H*, our investor could transfer consumption from low-value situations to high-value situations and make himself better off.
- Consumption in situations *LG* and *LB* ($40,000 on average) is higher (and thus less valuable) than in situation *L* ($30,000). By saving less in situation *L*, our investor could transfer consumption from low-value situations to high-value situations and make himself better off.

Now, suppose we introduce health insurance. And suppose health insurers offered a contract that says that for every dollar you invest, you receive $2 if your health is bad in retirement. If you borrow a dollar in prime earning years and invest in this stylized form of health insurance, what happens? Your consumption in prime earning years remains the same but falls by $1 when your health is good (because you have to repay the loan) and increases by $1 when your health is bad (because you receive the $2 for having bad health less the loan repayment). Thus, health insurance transfers consumption resources from situations where your health is good

[3] Moving consumption over the life cycle is at the heart of life-cycle planning. As Irving Fisher put it, the intent of life-cycle planning theory is to "modify [the income stream received by the individual] by exchange so as to convert it into that particular form most wanted by [this individual]" (Fisher 1930, Chapter 6).

to situations where your health is bad. For example, in Figure 1, consumption is higher (and thus less valuable) in good health situations than in bad ones, so our investor can make himself better off by buying insurance.

Five Key Concepts

With the overview of the life-cycle model complete, we now highlight five of the most important insights produced by the theory of life-cycle planning: (1) the notion of a lifetime budget constraint, (2) the relevance of contingent claims in life-cycle planning, (3) the trade-off imposed by varying costs of consumption over one's lifetime, (4) the role of risky assets, and (5) the asset allocation decision over the life cycle.

Insight 1: The Lifetime Budget Constraint. One of the great early insights in financial planning follows directly from Principle 2 above: Under certain conditions, household consumption over the life cycle depends entirely on the present discounted value of lifetime income and not on the evolution of income itself. More specifically, suppose two investors have some financial wealth and also expect some stream of labor income over their remaining working life.[4] Now, calculate the discounted present value of their future income, which we call "human wealth," and add it to their savings and call the sum "total wealth." According to the life-cycle model, under certain conditions, if two investors have the same total wealth, then their consumption decisions over the life cycle will be the same, regardless of the shape of their actual income profile.

To understand why we can ignore the profile of income across dates and random outcomes, return to Principle 2, which says that we can use financial assets to transfer consumption from one situation to another. A loan is a financial asset that allows you to increase consumption today in exchange for reducing consumption by the amount of the loan plus interest at a future date. What is the maximum amount an investor can consume today? It is her current savings plus the maximum amount she can borrow. What is the maximum amount she can borrow? It is, theoretically, the present discounted value of her future labor income. Thus, total wealth, as defined in the previous paragraph, measures the maximum amount an investor can transfer to the present. Now that our investor has transferred everything to the present, she can decide when to spend it, and using the same technology, she can transfer her wealth to the situations where she wants to consume. It is important to stress that the idea of transferring all lifetime income to the present is purely a hypothetical construct used as a way to measure lifetime resources using a single metric.

---

[4]The lifetime budget constraint is presented in Fisher (1930), Modigliani and Brumberg (1954), and Modigliani (1986). This concept of a lifetime budget constraint has been generalized and successfully applied to life-cycle planning under uncertainty, starting most notably with Cox and Huang (1989), and has been central to the development of finance theory over the past decade or so.

©2007, 2008, The Research Foundation of CFA Institute

The importance of the lifetime budget constraint is that it shows that financial wealth is only one part of the wealth that matters to an investor. Total wealth equals the sum of both financial wealth and human wealth. For most households—basically for almost all households in which the income earner is not old—human wealth dwarfs financial wealth. **Table 1** shows the ratio of human wealth to income measured using real-world data for various groups in the population. To see the importance of human wealth, consider a 35-year-old male college graduate earning $100,000 and owning $400,000 in financial wealth. Consider also an heir who plans to remain out of the labor force his entire life and has $3 million in financial wealth. One might think that these two investors have nothing in common, but according to one version of the life-cycle model, they should, in fact, consume exactly the same amounts. Note that according to Table 1, the college graduate's human wealth is 25.9 times his current income, or in this case, $2.59 million, to which we add financial wealth of $400,000 to get total wealth of $2.99 million, almost exactly the same as the heir.

We can also incorporate future expenses into the lifetime budget constraint. For example, suppose an investor knows that she will send two kids to college at given dates in the future. If we know how much that education will cost, we can simply subtract the present value of future education costs from current total wealth, just as we added future income.[5]

Insight 2: The Importance of Constructing "Contingent Claims." Previously, we argued that investors can use financial assets to transform their income and expense streams into the equivalent of financial wealth, but we were unspecific as to how. Now, we focus on how investors can actually affect these transformations. First, the easy case: If future income and expenses are certain, then investors can transform them into additions and subtractions from current wealth by simply borrowing and/or saving the appropriate amount. For example, if one knows that one will earn $100,000 five years from now and the interest rate is 5 percent, one can raise current liquid financial wealth by borrowing $78,350 and paying back the money plus $21,650 in interest from future earnings.

But in practice, things are not so easy. The main problem with calculating the lifetime budget constraint is random outcomes. To see why random outcomes are a problem, return to Figure 1. Following the logic above, our investor can convert his future labor income into current financial wealth by borrowing. But how much should he borrow? Along the "high-income" path, he earns $100,000 a year. So, if we assume an interest rate of zero, he could borrow $100,000 today and pay it back, say, at age 45. But suppose he does not get the high-income outcome but, instead,

[5] See, for instance, CollegeSure savings funds.

Table 1. Human Wealth Measured as a Fraction of Current Income

	Initial Age				
Item	25	35	45	55	65
High school graduates					
Men	29.7	19.1	12.8	8.2	—
	(= 718,530/24,199)	(= 629,378/33,005)	(= 439,494/34,301)	(= 219,269/26,814)	—
Women	27.9	16.6	11.9	8.6	—
	(= 379,592/13,606)	(= 317,191/19,159)	(= 202,351/16,997)	(= 101,256/11,784)	—
College graduates					
Men	47.4	25.9	15.9	8.7	—
	(= 1,483,412/31,297)	(= 1,483,295/57,264)	(= 1,212,542/76,385)	(= 691,057/79,566)	—
Women	32.9	20.1	13.3	7.0	—
	(= 881,762/26,808)	(= 792,354/39,424)	(= 580,133/43,506)	(= 266,430/38,064)	—
Male college graduates					
Retire at 55	39.3	20.0	9.9	—	—
	(= 1,231,486/31,297)	(= 1,144,728/57,264)	(= 757,535/76,385)	—	—
Retire at 75	51.0	28.6	18.6	12.2	7.0
	(= 1,597,261/31,297)	(= 1,636,299/57,264)	(= 1,418,165/76,385)	(= 967,398/79,566)	(= 432,870/61,491)
Male advanced degree holders					
	51.0	28.1	16.7	9.2	—
	(= 1,651,729/32,386)	(= 1,709,956/60,773)	(= 1,431,000/85,798)	(= 866,733/94,627)	—

Note: Human wealth is measured as a multiple of current income. For example, a 25-year-old male college graduate has human wealth equal to 47.4 times his current income.

Source: Authors' calculations based on data from the Panel Study of Income Dynamics.

earns "low income." Then, at age 45, his income would not be sufficient to pay off the loan. Another alternative would be to borrow $75,000, the average of his two income draws, but he would still have insufficient funds.

What if financial markets offered a security that paid off $1 only if our investor got the low-income outcome and another security that paid off $1 only if our investor got the high-income outcome? Then, our investor could convert his future income along the low-income path into current income by shorting $50,000 of the low-income asset and similarly convert his income along the high-income path using the high-income asset. We call these assets that pay off contingent on some future event "contingent claims."[6]

If contingent claims are so useful, one might ask why we do not observe them. And the answer is that we do, although they rarely ever appear in the form we described. Let us revisit the case of health insurance. Earlier, we described a stylized contract that paid off only in bad health states. In its most familiar form, however, health insurance is a contract that pays the investor's medical bills both in good and bad health states in exchange for a payment today. Accordingly, health insurance—as we observe it—is not a contingent claim per se because it pays off in both outcomes. Nonetheless, it is easy to construct a contingent claim using this common form of health insurance and the riskless bond. For example, if you want a claim contingent on the bad health outcome, then borrow $5,000 and buy health insurance. In the bad health event, you receive $10,000 ($15,000 less the loan repayment), and in the good health event, you get nothing because the health insurance payoff and the loan repayment cancel each other out. One can similarly construct a claim contingent on the good outcome.[7]

Contingent claims help with another serious problem: inflation. Suppose our investor knows for sure that he will earn $100,000 next year, but he is uncertain about inflation. Suppose that inflation could either be zero or 10 percent. Thus, our investor's real income next year actually does vary randomly: Along one path, he receives $100,000 in real spending power, and along the other, he receives $90,000. According to Principle 2, our investor will want to shift consumption from the low-inflation event to the high-inflation event. If we create inflation-contingent claims, our investor will be able to do just that. It is for this reason that economists have long advocated and spearheaded the creation of inflation-indexed bonds, marketed as TIPS or Treasury Inflation-Protected Securities.[8]

---

[6] Also called Arrow–Debreu securities after the seminal work of Arrow (1953) and Debreu (1959). See also Arrow (1971) and the recent work by Sharpe (2006).

[7] This type of reasoning is the essence of production technology supporting the creation of contingent claims, in theory and in practice, going back to Black and Scholes (1973), Merton (1973), and Cox, Ross, and Rubinstein (1979). This is the core of modern financial engineering as it applies to life-cycle planning.

[8] See, for instance, Fischer (1975) and Bodie (2003) for the role of inflation-protected bonds in life-cycle plans.

We have now, we hope, convinced the reader of the value of contingent claims in life-cycle financial planning. But before we go on, we should discuss the limits to the use of contingent claims. First, contingent claims work well when both parties can verify the event in question and neither party can affect or has better information about the likelihood of the event occurring. For example, it is easy to verify the price of General Motors Corporation (GM) stock. And for the most part, investors who either have better information or who can affect the price are legally forbidden from trading, so we see a large variety of claims contingent on the level of GM stock (futures, options, etc.). But earlier, we proposed that an investor would buy claims contingent on the level of his labor income. First, because income is not always easy to verify (the investor would have an incentive to hide some income and claim that he earned only $50,000 when he actually earned $100,000 so as to allow him to pay back only $50,000) and because a worker has some control over how much he earns (again, our investor has an incentive to earn less because he would then have to repay less), income-linked contingent claims present practical problems.[9] Second, the creation of contingent claims requires that we clearly understand the risks investors face, and neither record keeping nor econometric techniques have yet rendered measurement of these risks trivial.

Insight 3: The Prices of Securities Matter! In discussing contingent claims, we have shown how households can eliminate variation in consumption across different random events by transferring consumption from outcomes with high consumption to outcomes with low consumption by buying and selling consumption in those different outcomes using contingent claims. But we have said nothing about the price of contingent claims, and as everyone knows, prices usually play a central role in determining how much of something someone wants to buy or sell.

To illustrate some of the issues with the pricing of contingent claims, we consider an investor who faces two equally probable outcomes in the future: In outcome *H*, she consumes $100,000, and in outcome *L*, she consumes $50,000. **Table 2** provides information for this example. To analyze this problem, we will need two concepts from probability theory. First, we measure the "expected" level of consumption, which we get by weighting different outcomes by their probability and summing. The expected consumption for our investor is $75,000. But many different consumption profiles yield the same expected consumption (for example, $75,000 with certainty), and risk-averse investors are not indifferent among them. Our investor, if risk averse, would certainly prefer $75,000 with certainty. But we can go further and actually measure how much she prefers other consumption

[9]These problems are related to moral hazard and the resulting borrowing constraint literature, such as He and Pages (1993) and El Karoui and Jeanblanc-Picque (1998), who extended the study of human capital in life-cycle planning to address the important real-world problems caused by the limited ability of people to borrow against future income.

©2007, 2008, The Research Foundation of CFA Institute

Table 2. Understanding the Role of the Prices of Contingent Claims

Item	H		L
Scenario			
Probability	50%		50%
Initial consumption	$100,000		$50,000
Expected consumption		$75,000	
Certainty equivalent		70,710	
Baseline prices			
Price of contingent claim	50¢		50¢
Strategy	Sell $25,000		Buy $25,000
Cost	–$12,500		$12,500
New consumption	75,000		75,000
Expected consumption		$75,000	
Certainty equivalent		75,000	
Alternative prices			
Price of contingent claim	40¢		60¢
Strategy 1	Sell $25,000		Buy $25,000
Cost	–$10,000		$15,000
New consumption	75,000		75,000
Expected consumption		$75,000	
Certainty equivalent		75,000	
Strategy 2	Sell $30,000		Buy $20,000
Cost	–$12,000		$12,000
New consumption	70,000		70,000
Expected consumption		$70,000	
Certainty equivalent		70,000	
Strategy 3	Sell $12,500		Buy $8,333
Cost	–$5,000		$5,000
New consumption	87,500		58,333
Expected consumption		$70,833	
Certainty equivalent		71,440	

profiles by measuring the "certainty equivalent consumption level"—the level of certain consumption that would make her as happy as the random consumption in question. For a particular level of risk aversion, we calculate that our investor would be as happy with $70,710 with certainty as with the $100,000 and $50,000 with equal probability.[10]

[10]This type of calculation, known as certainty equivalent analysis, is explained at http://en.wikipedia.org/wiki/Certainty_equivalent. Certainty equivalent measures have been used to compare investment strategies since at least Merton and Samuelson (1974).

Now, a financial planner comes along to help her out. The financial planner has at his disposal a set of contingent claims, one paying \$1 in outcome *H* and another paying \$1 in outcome *L*. What strategy should he propose to the investor? Let us start with a baseline case where both contingent claims cost the same: 50 cents. The financial planner proposes to the investor that she short \$25,000 of outcome *H* income by shorting the contingent claim and that she go long \$25,000 of outcome *L* income by purchasing the outcome *L* contingent claim. The cost of the state *L* contingent claims (\$12,500) is exactly offset by the gains from shorting the outcome *L* contingent claim, meaning that the long–short portfolio costs nothing. What happens to consumption? The investor now consumes \$75,000 in both outcomes. This strategy, therefore, shifts consumption from the high outcome to the low outcome—just what we said financial assets were supposed to do. And the certainty equivalent level of consumption, trivially equal to the actual level of consumption, \$75,000, is much higher than the initial level. So, the financial planner provided good advice.

Now, we change the world a little and set the prices of the contingent claims unequally at 40 cents for the *H* outcome and 60 cents for the *L* outcome. Suppose the planner provides the same advice (called "Strategy 1" in the table). The plan still provides the same certain level of consumption, but a slight problem exists: Because the price of the *H* state consumption has fallen relative to the *L* state, the revenue from the sale of *H* state consumption no longer offsets the cost of the added *L* state consumption. The investor now needs to come up with \$5,000 to execute the strategy. We assume that the investor does not have that money and confine ourselves to self-financing strategies. The planner regroups and provides Strategy 2, a self-financing strategy that yields certain consumption of \$70,000. By selling more of the cheap state *H* claims and buying fewer of the expensive state *L* claims, our investor can achieve certain consumption of \$70,000 without adding any money. Has the planner earned his money? No. Recall that certainty equivalent consumption for the initial consumption profile exceeded \$70,000.

Not all investors will respond to asset prices in the same way. Differences in risk aversion, for example, play a big role. Risk aversion measures the willingness of an investor to tolerate variation in consumption across random outcomes. In Strategies 2 and 3, we argue that our investor is willing to accept a substantial increase in variation of consumption in exchange for a small increase in expected consumption. But that conclusion follows only because of the specific choice of risk aversion that we made. A household with higher risk aversion might opt for the sure consumption. Another issue that has drawn significant attention from economists is "habits." Some have argued that households are particularly unwilling to tolerate reductions in consumption. In this case, for example, suppose that our investor currently consumes \$70,000 a year and is unwilling to tolerate any reduction in consumption. Then, for her, the only plausible option would be to accept the

©2007, 2008, The Research Foundation of CFA Institute

lower expected level of consumption that accompanies the strategy of achieving riskless consumption. To add to the challenge of portfolio choice in this situation, we point out that the investor knows that higher consumption may restrict her choices in the future and, therefore, may restrict consumption now so that she can take risky bets in the future.[11]

So far, both strategies proposed by the planner have caused problems—one because it required substantial additional funds and the other because it did not provide any benefit. Does this mean that financial planning cannot help this investor? No. Strategy 3 offers a bundle of contingent claims that manages to raise the investor's certainty equivalent consumption without requiring additional investment. What is unique about this plan among the ones we have looked at so far is that it does not provide a certain level of consumption.

The earlier examples illustrate that the optimal plan depends on the prices of the contingent claims. With either set of prices, it was possible to eliminate variability from consumption. But in only one of the cases was this advisable. The difference between the two scenarios is that in the baseline, there was no risk premium (and thus no incentive for the investor to take on risk) and in the alternative, there was a risk premium. To see the difference, calculate the returns on a riskless bond that pays \$1 in both outcomes and an equity-type security that pays \$1 in state *L* and \$2 in state *H*. In the baseline scenario, the price of this bond would be \$1, making the return zero; the price of this equity would be \$1.50 with an expected payoff of \$1.50, so it also has a return of zero. In the alternative scenario, the return on the riskless bond is still zero, whereas the equity security now costs \$1.40 with the same expected payoff of \$1.50, meaning that it returns 7 percent more than the riskless bond.

Insight 4: Risky Assets in the Life-Cycle Model. One of the most important insights of the life-cycle model concerns the benefits of risky assets. In the life-cycle model, we view risky assets as a way to move money across different outcomes at a given time, not as a way to transfer resources across time. Let us illustrate this point with an example.

Consider an investor who lives for two years, this year and a "next year" in which there are two possible outcomes: "good times" and "bad times." Our investor can invest in a bond that pays 5 percent regardless of the outcome and a stock that increases 30 percent in good times and falls 5 percent in bad times. **Figure 2** illustrates this event tree. Standard investment advice would view the two assets as different ways to save for the future—to transfer money to the future. In the life-cycle model, we divide their roles. We do so by constructing a portfolio composed

[11] For a discussion of these issues, see Dybvig (1995).

Figure 2. The Role of Risky Assets in the Life-Cycle Model

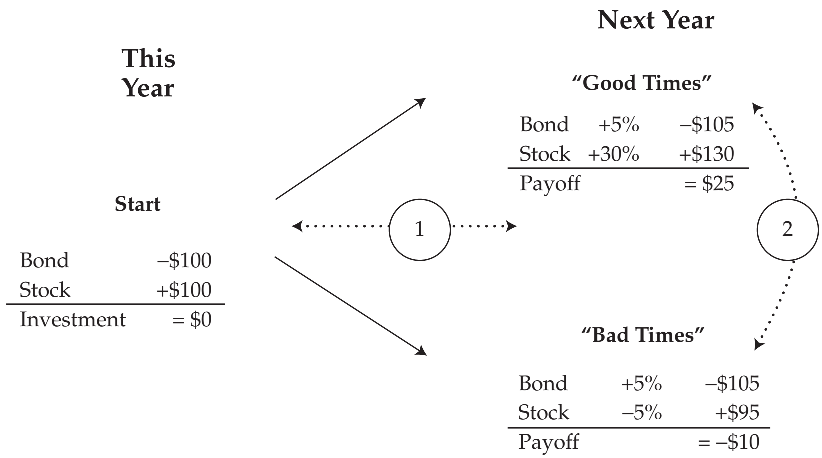

of $100 of stock financed by a $100 short position in the bond. This portfolio costs nothing today and, as Figure 2 shows, pays $25 in good times and –$10 in bad times. In other words, one can use this portfolio to convert $10 in bad times into $25 in good times (circled Transaction 2 in the figure). To transfer money across time, one uses the bond that, for example, allows the investor to transfer $100 today for $105 in both states in the future (circled Transaction 1 in the figure).

Suppose our investor wants to decide whether to invest $100 in the stock or bond. Our analysis shows that a $100 investment in the stock is a combination of a $100 investment in the bond and the portfolio described above. In other words, the investor is exchanging $100 today for $105 in the future and exchanging $10 in bad times for $25 in good times. According to our logic, we view investing $100 in the bond as exchanging $100 today for $105 next year but transferring nothing from bad times to good times. Thus, the difference in the two investment options has to do not with transferring resources across time—both investments achieve that—but with transferring resources across outcomes.

Therefore, the decision to invest in the stock revolves around whether the investor is willing to give up $10 in bad times in exchange for $25 in good times. If we imagine, for the purposes of discussion, that the two outcomes are equally probable, then this seems like an exceptionally good deal. But we should recall that a goal of financial planning is to smooth consumption across outcomes, so we need to know whether our investor wants to transfer income from bad times to good times. At first glance, one might imagine that our investor would want to transfer income just the opposite way if, for example, "good times" meant employment and "bad times" meant joblessness.

©2007, 2008, The Research Foundation of CFA Institute

Insight 5: Asset Allocation over the Life Cycle. One of the great early discoveries of the theory of life-cycle financial planning was an understanding of the evolution of the optimal level of risk exposure as an investor ages. Despite the prevailing folk wisdom (and advice from some practitioners and academics) that investors should reduce the proportion of their portfolio invested in risky assets as they age, Merton (1969) and Samuelson (1969) showed that, for standard preferences,[12] it is optimal for individuals to maintain a constant fraction of their total wealth (human capital plus financial wealth) invested in equities at all ages.

The above may suggest that the life-cycle theory has little advice on asset allocation other than to choose the right proportion once. In fact, because of the influence of labor income, the proportion of financial wealth invested in risky assets can vary dramatically over the life cycle. This issue was taken up by Bodie, Merton, and Samuelson (1992), who considered life-cycle investors with risky wages and a degree of choice with respect to the labor–leisure decision. The model's results indicate that the fraction of an individual's financial wealth optimally invested in equity should "normally" decline with age for two reasons. First, because human capital is usually less risky than equity and because the value of human capital usually declines as a proportion of an individual's total wealth as one ages, an individual may need to invest a large share of his or her financial wealth in risky assets to get sufficient overall risk exposures. Second, the flexibility that younger individuals have to alter their labor supply allows them to invest more heavily in risky assets. The opposite result, however, is also possible. For people with risky human capital, such as Samuelson's businessmen or stock analysts, the optimal path may be to start out early in life with no stock market exposure in one's investment portfolio and increase that exposure as one ages.

New Financial Products

In this section, we map the key concepts from the life-cycle model into a methodological framework for the successful production of consumer-oriented structured life-cycle contracts. This approach, which recognizes that individuals do not have the ability to engineer for themselves contingent claims packaged in the form they most desire, owes much to Merton's development of the theory of financial intermediation.[13] The production of life-cycle products is successful if

1. the state-contingent consumption promised by the life-cycle product effectively satisfies the needs of the individual and
2. the financial intermediary has the adequate production/risk management infrastructure in place for the consumer-oriented life-cycle product to be technologically feasible.

[12]In particular, Samuelson (1969) and Merton (1969) studied individuals with constant relative risk aversion. For more details, see http://en.wikipedia.org/wiki/Risk_averse#Relative_risk_aversion.

[13]This approach is exposed in Chapter 14 of Merton (1992).

An analogy is that of the automobile industry. The customer is offered a standard product with optional features. Then, it is the automobile manufacturer's responsibility to fill the order received. Similarly, the development of financial products should start with designing standard products that serve primarily to secure a certain standard of living in retirement with various options.

In particular, the technology applied to this approach is entirely nested within the contingent claims analysis (CCA) framework developed by Merton (e.g., 1992). In accordance with the CCA methodology, the process that allows a financial intermediary to successfully engineer, market, and deliver optimal life-cycle products to its customers is best viewed as a production process that consists of three sets of tasks, as illustrated in **Figure 3**.

Figure 3. Using the Life-Cycle Model to Address Consumer Needs

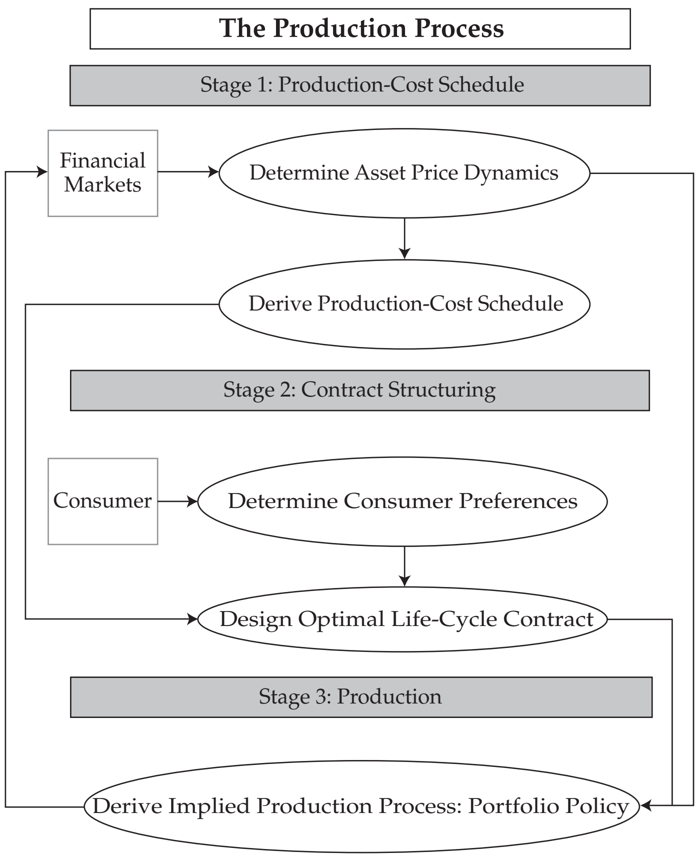

Stage 1. This stage requires the financial firm to determine its production-cost schedule (i.e., the cost of state-contingent consumption). These costs fall under two categories: the cost of certain future consumption provided by the market prices of inflation-protected bonds and the cost of state-specific consumption implied by the market prices of such securities as options (e.g., Breeden and Litzenberger 1978). State-contingent prices, also known as Arrow–Debreu prices (Arrow 1953; Debreu 1959), allow one to gauge the objective trade-off of consuming in different states of the world; thus, they are critical inputs in the process of designing optimal life-cycle consumption programs. These prices are instrumental in guaranteeing that the life-cycle contract satisfies the individual's budget constraint. A critical aspect of this stage is that it requires no knowledge of the financial intermediary's future customers and their preferences.

Stage 2. This is a contract-structuring stage during which, using the schedule of state prices, the financial firm structures the product that best satisfies the customer's tastes given the individual's lifetime wealth and planning horizon. For anticipated spending items with a well-defined time structure, and for which successful funding is a high priority—such as securing a comfortable standard of living in retirement and being in a position to send children to college—the financial firm should guarantee a floor given a particular flow of contributions. To the extent that the individual has wealth in excess of that required to finance necessities, modern finance theory (e.g., Vasicek 1977; Cox and Huang 1989; and Merton 1992) indicates that optimal discretionary spending over the life cycle is functionally identical to a derivatives contract on the growth-optimum portfolio.[14] By acknowledging the effects of consumption habit formation (e.g., Dybvig 1995; Bodie, Detemple, Otruba, and Walter 2004) on individual well being, "consumption ratcheting" features may play a central role in the design of consumer-oriented life-cycle products.

Stage 3. After the contract is finalized, the financial intermediary initiates production by dynamically replicating the payoffs promised to the customer. This is done via hedging in financial markets.[15] Although the financial intermediary replicates the guaranteed payments as closely as possible, it must also raise risk capital from investors to assure customers that their payments will not be subject to residual risk (e.g., the risk that the intermediary may default on its obligation to deliver the promised amounts).[16] Consumers should not have to become investors in the debt or equity of the financial intermediary.

---

[14]The growth-optimum portfolio minimizes the expected time to reach a prespecified wealth target under the assumption of continuously reinvested gains. The composition of the growth-optimum portfolio is subjective in that it depends on expected rates of return on stocks.

[15]For example, Detemple, Garcia, and Rindisbacher (2003, 2005) offer a flexible mathematical technology to determine the precise portfolio policy supporting optimal life-cycle products.

[16]Risk capital is defined in Merton and Perold (1993).

Conclusion

As products and services for addressing the financial risks of retirement change, so do the varieties of institutions available to provide support to the elderly. Today, many diverse retirement income systems coexist around the world, each relying in varying proportions on one or more of the following institutional forms:

- support from family or community,
- pension plans sponsored by employers and/or labor unions,
- social insurance programs run by governments, and
- personal savings in the form of real and financial assets—equity in one's home or business, savings accounts, insurance contracts, mutual funds, and so on.

Many experts agree, however, that the mix of these institutional forms will change significantly in the next few years. Change is particularly to be expected in the industrialized countries, such as the United States, the United Kingdom, Australia, Western Europe, and Japan, where the rapid aging of the population reflects both that people are living longer and that they are having fewer children. In these economies, people will find they can rely less on family and government support than in the past and must, instead, turn to financial markets and related institutions by saving and investing for their own retirement. Even in emerging markets, new demographic and economic realities have prompted the beginning of widespread retirement system reforms, as seen in the pension reform movements of Latin America, Eastern Europe, and more recently, Asia.

In response to global population aging and financial deregulation, governments and financial firms are seeking to create new institutions and services that will provide the desired protection against the financial consequences of old age, illness, and disability and will insulate people against both inflation and asset price fluctuations. New opportunities are to be expected for older persons to continue employment, perhaps on a part-time basis, and to convert their assets, particularly housing wealth, into spendable income. For better or for worse, these developments mean that people are being given more individual choice over their own asset accumulation and drawdown processes. Because these new financial instruments transfer more responsibility and choice to workers and retirees, the challenge is to frame risk–reward trade-offs and to cast financial decision making in a format that ordinary people can understand and implement.

The authors thank Phil Dybvig, Debbie Lucas, and participants at this conference for helpful comments and suggestions.

©2007, 2008, The Research Foundation of CFA Institute

REFERENCES

Arrow, K. 1953. "The Role of Securities in the Optimal Allocation of Risk-Bearing." *Econometrie* [translated and reprinted in *Review of Economic Studies* (1964), vol. 31, no. 2:91–96].

———. 1971. "Insurance, Risk, and Resource Allocation," Chapter 5 in *Essays in the Theory of Risk-Bearing*. Chicago: Markham Publishing Company.

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

Bodie, Z. 2003. "Thoughts on the Future: Life-Cycle Investing in Theory and Practice." *Financial Analysts Journal*, vol. 59, no. 1 (January/February):24–29.

Bodie, Z., R. Merton, and W. Samuelson. 1992. "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model." *Journal of Economic Dynamics and Control*, vol. 16:427–449.

Bodie, Z., J. Detemple, S. Otruba, and S. Walter. 2004. "Optimal Consumption-Portfolio Choices and Retirement Planning." *Journal of Economic Dynamics and Control*, vol. 28:1115–1148.

Breeden, D., and R. Litzenberger. 1978. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business*, vol. 51, no. 4 (October):261–651.

Cox, J., and C. Huang. 1989. "Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process." *Journal of Economic Theory*, vol. 49, no. 1 (October):33–83.

Cox, J., S. Ross, and M. Rubinstein. 1979. "Option Pricing: A Simplified Approach." *Journal of Financial Economics*, vol. 7, no. 3 (September):229–263.

Debreu, G. 1959. *Theory of Value*. New Haven, CT: Yale University Press.

Detemple, J., R. Garcia, and M. Rindisbacher. 2003. "A Monte Carlo Method for Optimal Portfolios." *Journal of Finance*, vol. 58, no. 1 (February):401–446.

———. 2005. "Intertemporal Asset Allocation: A Comparison of Methods." *Journal of Banking and Finance*, vol. 29, no. 11 (November):2821–2848.

Dybvig, P. 1995. "Dusenberry's Ratcheting of Consumption: Optimal Dynamic Consumption and Investment Given Intolerance for Any Decline in Standard of Living." *Review of Economic Studies*, vol. 62, no. 2 (April):287–313.

El Karoui, N., and M. Jeanblanc-Picque. 1998. "Optimization of Consumption with Labor Income." *Finance and Stochastics*, vol. 2, no. 4:409–440.

Fischer, S. 1975. "The Demand for Index Bonds." *Journal of Political Economy*, vol. 83, no. 3 (June):509–534.

Fisher, I. 1930. *The Theory of Interest*. New York: Macmillan Company.

He, H., and H. Pages. 1993. "Labor Income, Borrowing Constraints, and Equilibrium Asset Prices." *Economic Theory*, vol. 3, no. 4:663–696.

Merton, R. 1969. "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case." *Review of Economics and Statistics*, vol. 51, no. 3 (August):247–257.

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

———. 1992. *Continuous-Time Finance*. Malden, MA: Blackwell.

Merton, R., and A. Perold. 1993. "Management of Risk Capital in Financial Firms." In Chapter 8 of *Financial Services: Perspectives and Challenges*. Edited by S.L. Hayes III. Cambridge, MA: Harvard Business School Press:215–245.

Merton, R., and P. Samuelson. 1974. "Fallacy of the Log-Normal Approximation to Optimal Portfolio Decision-Making over Many Periods." *Journal of Financial Economics*, vol. 1, no. 1 (May):67–94.

Modigliani, F. 1986. "Life Cycle, Individual Thrift, and the Wealth of Nations." *American Economic Review*, vol. 76, no. 3 (June):297–313.

Modigliani, F., and R. Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data." In *Post Keynesian Economics*. Edited by K. Kurihara. New Brunswick, NJ: Rutgers University Press.

———. 1979. "Utility Analysis and Aggregate Consumption Functions: An Attempt at Integration." In *Collected Papers of Franco Modigliani*, vol. 2, Edited by A. Abel. Cambridge, MA: MIT Press.

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

Sharpe, W. 2006. "Retirement Financial Planning: A State/Preference Approach." Working paper.

Vasicek, O. 1977. "An Equilibrium Characterization of the Term Structure." *Journal of Financial Economics*, vol. 5, no. 2 (November):177–188.

©2007, 2008, The Research Foundation of CFA Institute