
Investing for a Distant Goal: Optimal Asset Allocation and Attitudes toward Risk

In this issue:

Introduction

Making Judgments about Risk

Expected Utility: A Framework for Thinking about Choices in the Presence of Risk

The Certainty Equivalent: Quantifying the Cost of Risk

Dynamic Programming

Putting It All Together: A Detailed Example

Results of a Simulation Exercise

Costs of Not Making the Best Allocation Choice

Conclusions

A hallmark of a defined contribution pension plan is the opportunity it provides participants to shape their own retirement to reflect their individual needs and attitudes. At the same time, the opportunity presents a significant challenge, as participants are faced with the complex task of deciding how to allocate or re-allocate retirement plan investments best over time.

TIAA-CREF has always provided extensive education and information to participants in the form of pamphlets, booklets, seminars, software, and other media to assist them in making choices. Recently, TIAA-CREF has moved further, offering individuals specific asset allocation advice that considers their personal characteristics and financial situation, as well

as their comfort level regarding risk. The model underlying this advice was developed based on the extensive experience of TIAA-CREF staff, as well as on formal theoretical and empirical studies conducted by financial economists.

This article by Professor Jim Musumeci, Department of Finance, Southern Illinois University at Carbondale, summarizes his research on the subject of optimal asset allocation. It describes and presents an academic economist's approach to thinking about the issues involved in asset allocation. Professor Musumeci's work emphasizes the importance of individuals' attitudes toward risk and the changing level of their accumulations in determining the best allocations for those saving for a distant goal.

It is important to note that the opinions and analysis presented in this article are the results of the academic work of Professor Musumeci. The article is presented for your general interest only and should not be interpreted as or considered to be investment advice from TIAA-CREF.

Introduction

As the baby boom generation sees retirement looming on the horizon, we are witnessing an increased interest in the issue of asset allocation when saving for a distant goal. Books by prominent authors suggest that younger investors should invest aggressively, with a slight shift to more conservative investments only as the goal date nears. Similarly, academic research in this area suggests that an all-equity portfolio will earn a higher return than a fixed-income portfolio most of the time if the investor's goal is distant.

However, most of this research has ignored the investor's attitudes toward risk. Investors generally consider not just the average return but also what happens when returns are below average. If an investment strategy fails only infrequently, but fails miserably when it does, investors may reject it in favor of another strategy, which fails more often, but by smaller amounts.

Are equities the best investment vehicle for younger investors? Do optimal asset allocations depend on the risk attitudes of a specific investor or on how much that investor has saved to date? Are there generalizations we can make for most investors? How much does it "cost" an investor to choose an allocation that does not fit well with that investor's specific goals?

This paper examines some of these questions and finds that attitudes toward risk and progress toward achieving savings goals can have a significant effect on optimal asset allocations. Failure to consider such factors can be costly. This is not to say that the traditional "invest aggressively while you're young" advice is inappropriate for conservative young investors. If anything, the results suggest that the traditionally recommended asset allocations are probably *too conservative*. Performing even worse are portfolios of Treasury securities alone, even after their lower risk is taken into account. Finally, while an individual's attitudes toward risk affect optimal asset allocations, there are some generalizations that can be made for most investors.

Making Judgments about Risk

Suppose you were asked the following question: “Which would you prefer to earn: (a) \$50,000 next year for sure, or (b) \$100,000 next year with probability one-half and \$25,000 next year with probability one-half?”

In all likelihood this question would elicit different responses from different people. The second choice clearly produces greater income on average, but it also entails much greater risk. Whether you would choose the risky option or not depends on whether you feel the gain in expected salary more than compensates you for the increased risk. And while the average salary from the latter choice can be objectively calculated as $(1/2) \times \$100,000 + (1/2) \times \$25,000 = \$62,500$, the perceived risk is subjective and will vary from person to person. For some individuals, the gain in average salary of $\$62,500 - \$50,000 = \$12,500$ exceeds the (personal) “cost” of greater risk, and so they choose the riskier alternative. Others may find that “cost” of risk to be too great and so choose the certain \$50,000.

It does not surprise us that different consumers quite rationally choose to buy different products. Yet traditional asset-allocation advice often suggests that all investors should make the same choices. But in both the case of a consumer choosing among various products and that of an individual choosing from among various risky alternatives, the tradeoff between the costs and the benefits of making a particular choice is personal and may be different for different people. Virtually all we can say about any rational person is that, when presented with two alternatives having identical benefits but different costs, he or she will prefer the one with lower cost; and similarly, when facing alternatives with identical costs but different benefits, he or she will prefer the one with greater benefit.

The fact that different investors may quantify costs in different ways when dealing with investments makes it especially hard to make generalizations about what is the “best” choice for investors. A major difficulty is that we do not all have the same view of what “risk” is. While we

all may agree that risk is some measure of dispersion of returns, we do not agree on how to measure that dispersion. Is a small likelihood of great loss worse than a large probability of smaller loss? Does all dispersion around the average count, or only that below the average or below some target?

Expected Utility: A Framework for Thinking about Choices in the Presence of Risk

Economists have developed a modeling framework, called the theory of expected utility, that is intended to describe how individuals make decisions in the presence of risk. The theory is based on the concept of an “expected utility function.” An expected utility function is a mathematical, quantitative way of expressing the pleasure or satisfaction that an investor feels as a result of accumulating each successive dollar. It translates an individual’s attitudes toward wealth and risk into a mathematical object (a function) that can be used in calculations.

The expected utility function is useful as a theoretical construct because it completely describes an investor’s preferences with regard to both wealth and risk, and therefore provides a way to answer questions about how that individual assesses the costs and benefits of making particular choices. It is important to understand that the investor’s attitude toward risk is actually incorporated into the expected utility function. The framework of expected utility

also incorporates the important economic idea that while more wealth is better, every successive dollar buys something at least a little less important to the buyer than the previous dollar did.

All of this means that if we knew an investor’s expected utility function, and if we assume that the investor acts to maximize that function, then we could actually calculate which choice would be best for him or her when faced with any set of possible choices. (See the box below for a numerical example of how this works.)

In practice, however, most of us have no idea what our utility functions are. The problem is that to know a person’s utility function exactly would require much more information about that person than could ever be practically obtained, even from a very detailed questionnaire. As a result, implementation of the expected utility model is problematic. Such problems are typical of what Landsburg (1994) describes as “the artificial world of the economic model—a world in which everything is specified in the explicit detail that is never available in reality.”

This problem is an important one, but it certainly does not mean that the theory is useless. In particular, if a wide variety of different utility functions prescribe similar behavior, we can infer that it is reasonable for most investors to follow those prescriptions—even without knowing exactly what their individual utility functions are. Also, since the theory allows us

An Example of Expected Utility Calculations

Consider again the choice between receiving (a) \$50,000 next year for sure, or (b) \$100,000 next year with probability one-half and \$25,000 next year with probability one-half.

Suppose that your utility function is equal to the square root of your wealth. (Such a utility function is consistent with the preferences that a moderately risk-averse investor might have.) Your expected utility for each of the two alternatives, (a) or (b), is defined as the probability-weighted average of the functional value of the possible outcomes of each choice. For example, if you choose option (a), the only possible outcome is that you will receive \$50,000. Your expected utility for option (a) is therefore $1 \times (50,000^{.5})$, or 224. If instead you choose option (b), you might get \$25,000 or you might get \$100,000, each with probability one-half. The functional value (the square root) of the \$25,000 outcome is 158, and the value of \$100,000 is 316, so the expected utility of option (b) is therefore $(1/2) \times 158 + (1/2) \times 316 = 237$. The expected utility of the second alternative, (b), is greater than that of the first, (a), so (b) would be the better choice for you, based on this utility function. (Note that another investor with a different utility function might find (a) to be preferable. An example of such a utility function is $U[\text{wealth}] = \text{wealth} - \text{wealth}^2/400,000$.)

to account for risk in a very specific way, we could perform an analysis using a utility function that might be typical of an extremely conservative investor, and then use the prescriptions of that analysis as a way to judge the appropriateness of our own behavior. Because the expected utility framework is rigorously defined and logically consistent, it provides a solid basis for drawing conclusions. Other modeling frameworks may lack such strong foundations. For these reasons and others, a wide variety of academics from various disciplines agree that the expected utility model is the appropriate method to use when thinking about how to make the best choices in a risky environment.¹

The Certainty Equivalent: Quantifying the Cost of Risk

Another use of the utility function is that it can produce a “certainty equivalent” for any set of possible outcomes. For example, if you found yourself exactly indifferent between (a) \$50,000 next year for sure or (b) \$100,000 next year with probability one-half and \$25,000 next year with probability one-half, then we would say that your certainty equivalent for option (b) is \$50,000.

The certainty equivalent is useful for assessing a specific investor’s “cost” of risk, i.e., the (subjective) loss in value due to increased risk. We can also use it to find the cost of a choice that is suboptimal—i.e., not the best possible. For example, suppose your certainty equivalent of option (b) were only \$40,000, and without consulting you, I had to choose either (a) or (b) on your behalf. If I chose (b), we could describe my selection as having a \$10,000 “cost” to you, in the sense that I chose an option with a \$40,000 certainty equivalent when I could have chosen the certain \$50,000.

Dynamic Programming

The expected utility model described in the previous section is useful for finding the optimal allocations in a one-period framework, i.e., when one year remains before the investor’s goal date. However, accumulating a large targeted amount of money is often a long process, and it is not immediately clear how knowing what

asset-allocation decision to make one year before the goal date will help an investor make a similar decision, say, thirty years earlier. Adapting the engineering tool of dynamic programming allows us to determine optimal allocations at any time during a saver’s accumulation period.

A wide variety of academics from various disciplines agree that the expected utility model is the appropriate method to use when thinking about how to make the best choices in a risky environment.

Dynamic programming is perhaps best described by example. Imagine you were responsible for sending a spacecraft to a particular point on the moon. There are so many factors to consider—the earth’s rotation, the moon’s revolving around the earth, and various random events that are unknown at launch time—that it is difficult to account for them all at once in taking the appropriate first step.

A somewhat simpler approach is to break the problem down into a number of steps and work backwards from the desired end position. Thus the first problem to solve is actually the last step that will occur in real time, namely, how to land the spacecraft at the designated spot on the moon once it is in a lunar orbit. Once that problem has been solved, we can then direct our attention to the next-to-last step, how to get the spacecraft into a lunar orbit once it is 100 miles away from the moon’s surface. After solving this second problem, we can then consider how to get it 100 miles away when it is 200 miles away, etc. Eventually we will develop a sequence of instructions that tells us what to do at the first step, i.e., upon launch.

Putting It All Together: A Detailed Example

How can the tools of dynamic programming, certainty equivalents, and expected utility be used to find optimal asset allocations during a saver’s lifetime? The following example, depicted in Figure 1, gives a flavor of how optimal allocations can be determined.

Consider an investor who has accumu-

lated \$800,000 two years before reaching his goal date and is making no further contributions.² For simplicity, suppose there are only two asset classes, Treasury securities and equities, and two equally likely possibilities for financial market returns, either down or up. If the market is

down, then the return on Treasury securities is 8% and the return on equities is –10%. On the other hand, if the market is up, then the return to Treasury securities is 12% and the return on equities is 50%.³ Assume also that the investor considers only two possible allocations, (a) 100% to equities or (b) 0% to equities.

Finally, suppose the investor is concerned primarily with maximizing the level of wealth attained on a goal date, and that his attitude toward risk is such that when facing the same set of investment risks and returns, he would place the same total amount of wealth at risk, regardless of his overall level of wealth.

Two years before the goal is reached, the investor’s optimal allocation of the \$800,000 is not immediately clear. However, the number of possibilities is sufficiently limited that we can consider them all. Figure 1 displays the investor’s choices and the expected utility and certainty equivalent resulting from each of these choices.⁴

As in the “spaceship” analogy, the first thing to do is look at the last decision that must be made in real time: i.e., we start by determining an optimal allocation for each of the four levels of wealth that the investor could possibly have at the point in time one year before the goal date. For example, suppose the investor has accumulated savings of \$864,000 with one year to go before the goal date. (This would be the case if the investor had chosen to allocate 0% to equities in the previous year, and the market had gone down. This

Figure 1
Dynamic Programming Decision Tree
 (dollar amounts in thousands)

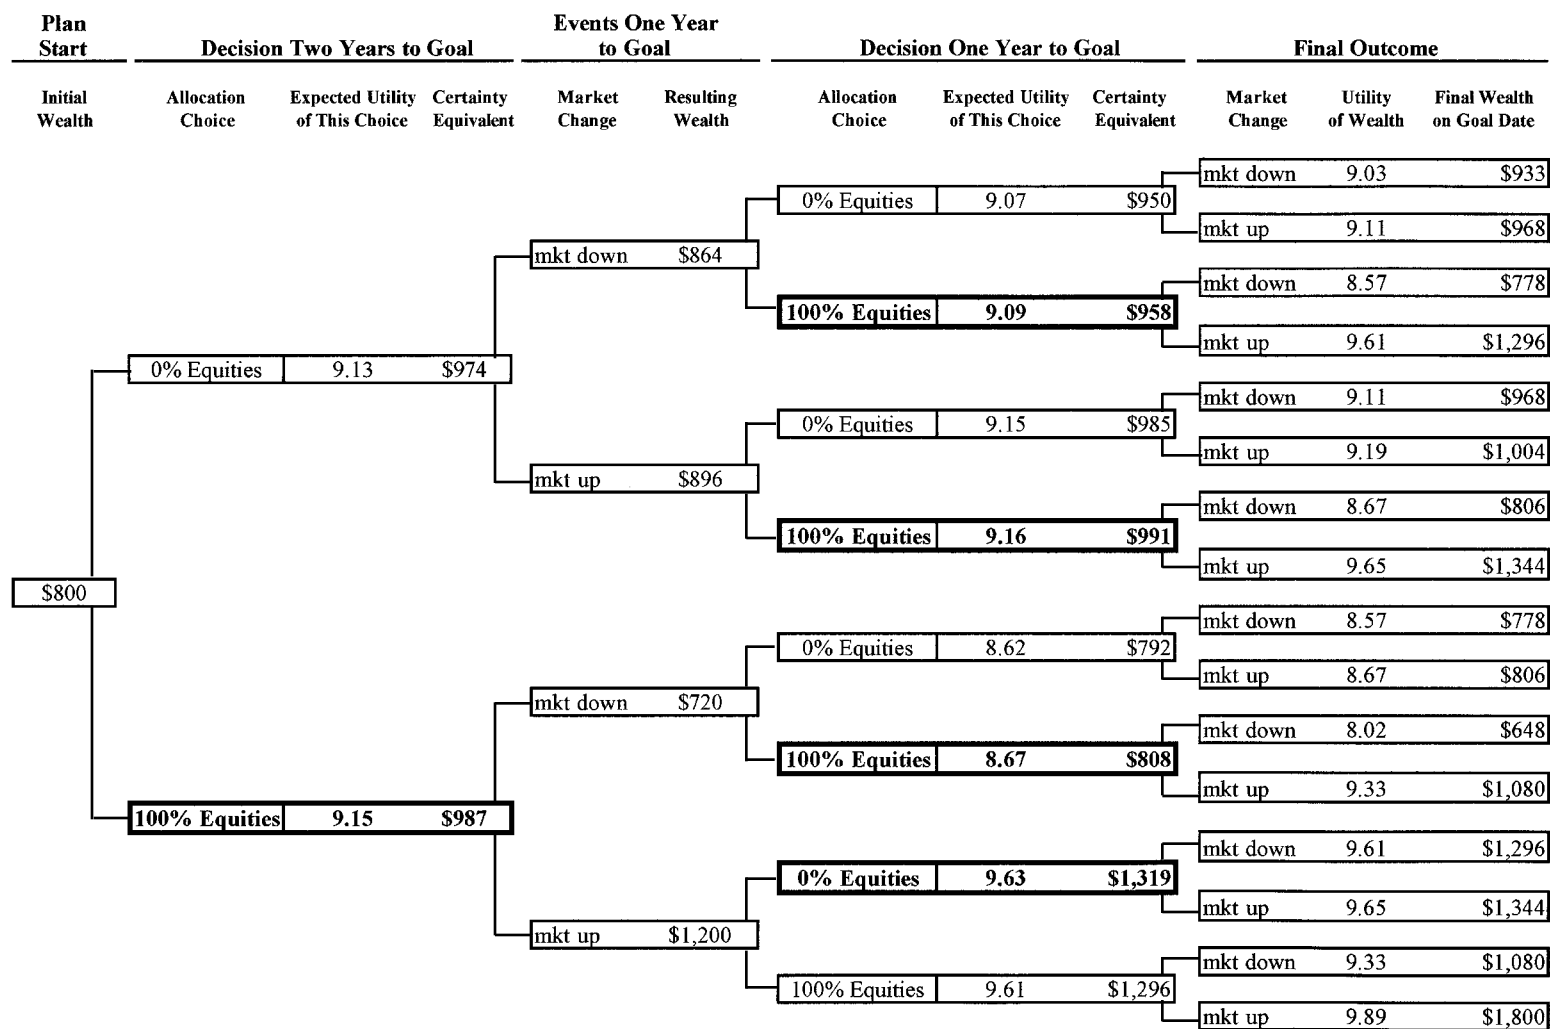


Figure 1 represents the results of the possible allocation strategies for each of the last two years before the goal date. The investor is assumed to begin with \$800,000 and to consider only two possible allocations in each period. For each possible level of wealth one year from the goal date, the investor can find the expected utility (or certainty equivalent) of each strategy and choose the higher one (boldfaced entries). Given the optimal allocations in the last period and their expected utilities, the investor can use the same procedure to find his or her optimal allocation to be 100% equities two years before the goal date.

sequence of events is depicted in the top “branch” of the decision tree in Figure 1.)

If at this point he chooses to allocate 0% to equities in the last year, the two possible outcomes are \$933,000 if the market goes down or \$968,000 if it goes up. The utility value of each of these two outcomes is shown to the left of each outcome in the last branch of Figure 1. Because these two outcomes are equally likely, the expected utility of making this allocation choice (0% to equities) is $(1/2) \times 9.03 + (1/2) \times 9.11 = 9.07$. This number has little intuitive meaning by itself, but it can be used to obtain the certainty equivalent, which, for the utility function used to generate the example, is \$950,000. For this investor, the value of choosing this allocation at this time is therefore equivalent to the value of receiving \$950,000 for certain at the goal date.

We can now compare the value of this allocation choice with the value of the alternative choice in order to find out which would be better for the investor at this point in time, with this level of accumulated wealth. The investor is still assumed to have \$864,000 with one year remaining until his goal, but now we look at the possible outcomes at the goal date when the investor chooses the 100% allocation to equities.

In this case, if the market goes down, he would end up with a utility of 8.57, while if the market goes up, he would have a utility of 9.61. This allocation choice therefore has an expected utility of 9.09 $[= (1/2) \times 8.57 + (1/2) \times 9.61]$, and a certainty equivalent of \$958,000. Since the expected utility (and certainty equivalent) for this choice is greater than for the choice of 0% to equities, we have determined that it would be optimal for the investor to allocate 100% to equities if he starts the last year with \$864,000. Also, we are now able to measure the cost, in terms of an “opportunity loss,” of making the wrong choice: Choosing the 0% allocation to equities would “cost” the investor the difference between \$958,000 and \$950,000, or \$8,000.

We can follow a similar process to determine optimal allocations for each of the other possible levels of wealth in the last

year before the investor’s goal date. It is important to note that the optimal allocation choice in the last year is sensitive to the level of accumulated wealth at that time. For example, suppose the investor begins the last year with \$1,200,000. (This would be the case if he had chosen the 100% equity allocation two years prior to the goal date, and the market had gone up. Such a chain of events is depicted in the bottom branch of the decision tree in Figure 1.) By comparing the appropriate expected utilities in this situation, we can see that the investor will now find an allocation choice of 0% to equities to be optimal, because its expected utility of 9.63 and certainty equivalent of \$1,319,000 exceed the 9.61 expected utility and \$1,296,000 certainty equivalent of the allocation of 100% to equities.

Once we have determined the optimal allocations and their expected utilities for all possible wealth levels in the last year before the goal date, we can use this information to determine optimal allocations for the next-to-last year as well, through a similar process. The process is in fact less complicated since we have assumed that there is only one possible level of wealth two years before the goal date (\$800,000). Given that the investor has \$800,000 at this point in time, we first figure out the value of choosing an allocation of 0% to equities. We do this by noting that if he makes this choice, he will then have either \$864,000 or \$896,000 in the last year before the goal date, depending on whether the market goes up or down. From our earlier calculations, we also now know the expected utility (and implied certainty equivalent) for the optimal allocation choices in the last period given the possible wealth levels at that time: 9.09 in the case of \$864,000, and 9.16 in the case of \$896,000.

Now, by assuming that the investor will in fact make these optimal allocation choices in the last year—regardless of whether the market goes up or down in the next-to-last year—we can calculate the expected utility of allocating 0% to equities in the next-to-last year: $(1/2) \times 9.09 + (1/2) \times 9.16 = 9.13$, implying a certainty equivalent of \$974,000. A simi-

lar calculation for the 100% equity allocation in the next-to-last year reveals that it has a higher expected utility of 9.15 (with a certainty equivalent of \$987,000). It is the assumption that the investor makes the best choice he can in every period that defines the optimal sequence of allocation choices throughout time: The method assures that at every point in time, the investor makes the best allocation choice he can, depending on how the uncertainty unfolds.

Thus we have determined that the investor should choose the allocation of 100% to equities in the next-to-last year before the goal date. We also know that the “cost” of making the wrong choice in the next-to-last year, *even when the optimal choice is made in the last year*, is $\$987,000 - \$974,000 = \$13,000$.

This example, although simplistic, suggests a tendency for the investor to become more conservative over time. With two years left until the goal, he allocates 100% to equities, but, given that allocation, with one year left he will allocate 0% to equities half the time and 100% to equities half the time.

Results of a Simulation Exercise

The detailed example in the previous section can be extended and expanded to determine optimal asset allocations when there are more asset choices and periods during which allocation decisions must be made. This section describes the results of a simulation exercise designed to determine the optimal asset allocation for investors saving for a distant goal date. (More details on exactly how the simulations were conducted can be found in the box on page 7.)

In principle, the dynamic programming technique can be used with any utility function. The simulations presented in this section are based on utility functions for three different hypothetical investors, each of whom is assumed to have a slightly different version of the same utility function. All three investors have a target level of wealth for a distant goal date. They all view risk in terms of the frequency with which wealth falls below the target wealth level, as well as the magnitude by which it does so.⁵

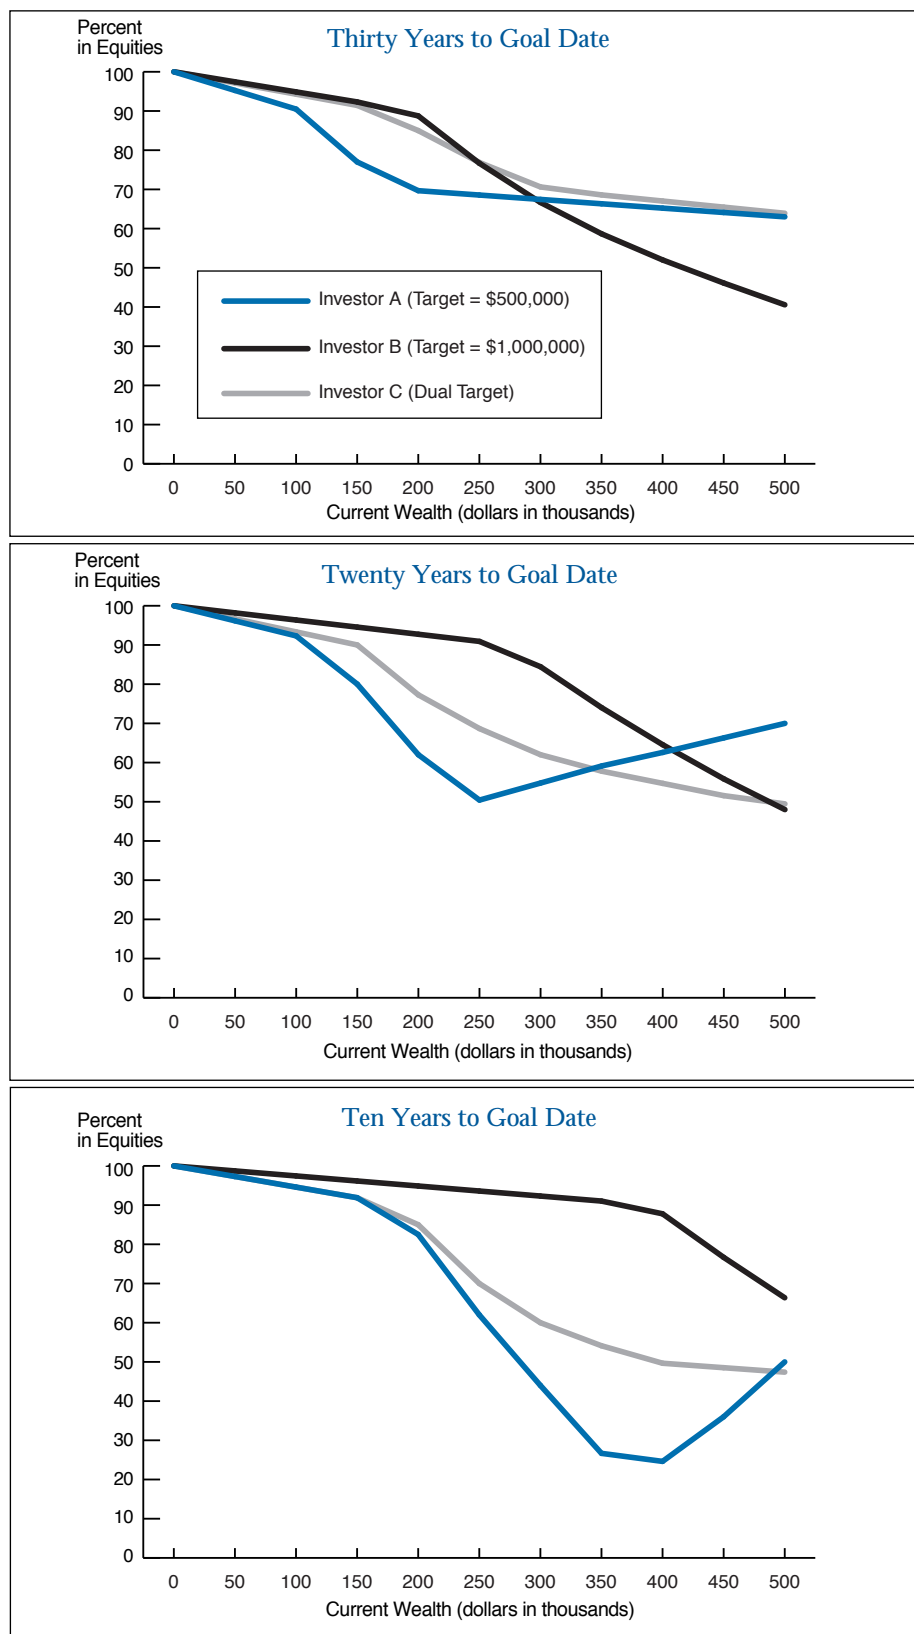
The investors differ from one another in terms of the target levels of wealth that they would like to have at their goal date. The first two investors, Investor A and Investor B, have a single target level of wealth in mind: Investor A has a wealth target of \$500,000 and Investor B has a wealth target of \$1 million.⁶ The other parameters of their utility functions are set in such a way as to be consistent with the preferences of highly risk-averse investors. In fact, these two investors are assumed to be so risk averse that, despite the much higher returns from equities and Treasury bonds, they will choose to invest all their wealth in Treasury bills if their accumulated wealth one year before retirement is at the target level.

The utility functions of these investors also imply that the dissatisfaction resulting from falling far below the target wealth level at the goal date would be dramatically more severe than the dissatisfaction resulting from a small shortfall. One of the characteristics of such a “shortfall risk” framework is that these investors become more cautious as their accumulated wealth approaches their targets. However, once the investors become fairly certain that they can achieve their target levels of wealth, they become less concerned about risk and once again invest more in equities. Because some very risk-averse investors might not be comfortable increasing their exposure to equities even after it’s very likely that their wealth target will be reached, we also model the behavior of a third investor, Investor C, who has a primary target wealth of \$500,000 and a secondary target wealth of \$1 million.⁷ Such a utility function assures that Investor C will maintain some aversion to risk, even when surpassing the first goal is very likely.

These particular utility functions were chosen in order to illustrate the robustness of the results we obtain and to provide benchmark predictions that would be appropriate for even highly risk-averse investors. The results of the simulation analysis using these three utility functions for these three investors are shown in Figure 2.

Perhaps the most remarkable feature of the simulation results is that, in all the

Figure 2
Optimal Allocations to Equities with Goal Dates Thirty, Twenty, and Ten Years in the Future



cases considered, an investor starting with any normal level of accumulated savings when the goal is forty years away will optimally invest all, or nearly all, of it in equities. (Because there is so little difference between the investors' optimal allocations forty years before the goal date, we do not chart these optimal allocations in Figure 2.)

Investors are still inclined to equities, although not as much so, when the goal is thirty years away, as depicted in the top panel of Figure 2. Here Investor A, with the \$500,000 target, switches from 100% equities to 90% equities if accumulated savings are over \$100,000. Investor B, with the target of \$1 million, will make the same switch at just under \$200,000. Not surprisingly, Investor C, the dual-target investor, behaves like Investor A at low levels of current wealth and like Investor B for higher levels of current wealth. This pattern of behavior is repeated again when fewer years remain until retirement. (In order to judge what kind of

accumulation might be typical at this point, note that if an investor started an investment program forty years before the goal date and were to make \$5,000 contributions at the end of every year, earning

simulation cases we do observe a reduced equity allocation twenty years before the goal date, equities remain the largest part of the portfolio at the levels of wealth that might typically be accumulated at that time.

The optimal allocations . . . vary greatly depending on the level of accumulated savings—a factor typically omitted from traditional allocation advice.

3% per year, he or she would have roughly \$57,000 when the goal was thirty years away.)

Only when the goal is about twenty years away do we begin to see significant divergence in the investors' optimal behavior. For example, in the middle panel of Figure 2, Investor A allocates only about 50% to equities if she has \$250,000 accumulated by this time, while Investor B allocates over 90% to equities with the same accumulated wealth. Although in some

For purposes of comparison, note that an investor starting an investment program forty years before the goal date, making \$5,000 contributions annually and earning 3% per year, would have roughly \$134,000 when the goal was twenty years away.

The middle panel of Figure 2 shows the first U-shaped allocations: Investor A, with the \$500,000 target, is increasingly cautious at accumulated wealth levels up to \$245,000. However, if her accumulated wealth at this time exceeds that level, she finds it optimal to invest greater amounts in equities. The reason that this behavior is optimal is that if her wealth at this time exceeds \$245,000, it is very likely that her target wealth will be either reached or at least nearly reached. Since the shortfall-risk approach assumes that risk becomes much less important once the target level of wealth is assured, the investor optimally allocates more to equities in order to earn higher average returns. Because Investors B and C are not as confident of reaching their goals at any of the wealth levels on the chart, their optimal allocation curves are not U-shaped in the ranges shown.

The bottom panel of Figure 2 depicts the optimal allocations for the investors when they are ten years from their goal date. This panel features the greatest divergence between the investors in terms of optimal allocations. For example, with \$400,000 in accumulated savings, Investor A allocates only about 25% to equities, while Investor B continues to allocate over 80% to equities.

The optimal allocations for these investors, however, vary greatly depending on the level of accumulated savings—a

Methodology for Simulations

The simulations used to derive the results described in the accompanying analysis are in principle the same as in the example of Figure 1. The main differences are:

- (1) the example considers only two possible allocations between two asset classes, while the simulations consider a wide variety of possible allocations among three asset classes (Treasury bills, Treasury bonds, and the S&P 500);
- (2) the example assumes no new contributions, while the simulations assume a constant real annual contribution of \$5,000; and
- (3) the example “rolls back” to only two years before the goal, while the simulations find optimal allocations for each year up to forty years before the goal.

The reported results are based on a simulation technique to produce the optimal asset allocations given the investor's attitudes toward risk, the number of years to the goal, and the accumulated savings to date. The asset classes considered are (1) Treasury bills, (2) long-term Treasury bonds, and (3) the S&P 500.

Ibbotson (1996) reports monthly and annual data for 1926-1995 and finds the average real equity return to be 9.2%; the average real T-bond rate to be 2.5%; and the average real T-bill rate to be 0.7%. We use the Ibbotson monthly data to simulate triplets of returns by choosing a year at random and taking the January return for each asset class, then repeating the process for February through December. The simulated annual return is the compounded value of the simulated monthly returns. For each year of the algorithm, 1,000 such triplets of annual returns are generated, and they are processed in the same way the two pairs of possible returns were in the example of Figure 1.

All models make simplifying assumptions, and it is important to note that our model looks only at wealth when the goal date is reached and assumes that is the end of the problem. This is a reasonable assumption if the investor's goal is the purchase of a life annuity. However, investors may also face allocation decisions when withdrawing money that are analogous to those they faced when accumulating it.

factor typically omitted from traditional allocation advice. For instance, while Investor A would allocate only about 25% to equities if she had \$400,000 accumulated savings, she would allocate over 90% to equities if her accumulated savings at this time were only \$150,000. In a pattern similar to what was observed in the middle panel of Figure 2, Investor A finds it optimal to invest more in equities if her accumulated wealth at this time is higher than \$370,000, since she is fairly confident of reaching her goal. Investors B and C are not assured of reaching their final target wealth level at any given level of current wealth on the bottom panel, so their allocation curves are not U-shaped. In order to judge what one might typically have accumulated after thirty years of planned saving, an investor starting an investment program forty years before the goal date and earning 3% per year would have roughly \$238,000 when the goal was ten years away.


It is natural to compare the three panels of Figure 2 with each other, but it is misleading to do so, since an investor should expect to have different amounts saved at the different points in time de-

picted in the charts. An investor with a target of \$1 million and accumulated savings of \$100,000 with a goal date forty years in the future is in a completely different situation from that of an identical investor with the same savings but a goal date only ten years in the future.

In order to illustrate what an optimal allocation pattern might look like over the entire forty-year accumulation period, Figure 3 depicts the optimal average equity allocation over time for an investor starting with no savings and making \$5,000 annual contributions. Also illustrated in Figure 3 are the asset allocation recommendations of some financial experts. As the figure suggests, the optimal allocations are, *on average*, fairly similar to those commonly suggested by the professionals. However, this averaging ignores the wide disparity in optimal allocations in later years. Allocations can vary greatly depending on the progress investors have made toward their savings targets. For example, the figure shows that the average allocation to equities in the thirty-ninth year was 41%, but in the 500 simulations this ranged from a low of 10% (in 162 cases) to a high of 100% (in 89 cases).

Finally, what if an investor is not optimistic that future returns on equities will be as large as past returns have been? We simulated optimal allocations for a fourth investor, Investor D, who is identical to Investor A of Figure 2 above (with a \$500,000 wealth target), but believes that monthly returns on equities are likely to be about a quarter of a percent lower in the future than they have been during 1926-1995. Our results indicate that even with this rather severe forecast of an approximate 3% drop in average annual equity returns, such an investor will allocate a large percentage to equities for small levels of current wealth if the goal date is forty years in the future. As shown in Figure 4, Investor D's optimal allocation to equities falls sharply when the goal date is only twenty or thirty years in the future, provided that significant progress has been made toward the target level of wealth (see the top two panels of Figure 4). For example, Investor A will allocate 70% to equities if she has accumulated savings of \$200,000 and the goal is twenty years away. On the other hand, Investor D, who is a less optimistic but otherwise identical investor, will allocate only 30% to equities under the same conditions.

Figure 3
Various Suggested Allocations to Equities



The two investors' optimal allocations converge, however, as the goal date nears. For example, Investor A will allocate 30% to equities if she has accumulated \$350,000 ten years before her goal, while Investor D will allocate 20% to equities.

Costs of Not Making the Best Allocation Choice

How expensive is it to make an allocation decision that is not optimal? One way of measuring this cost is in terms of the certainty equivalents discussed earlier. The certainty equivalent of an unknown amount is the certain value for which the investor would be willing to trade the unknown outcome. For example, if an investor is indifferent between receiving \$25 in cash or taking a chance on a coin flip that results in winning either \$60 or nothing at all, we would say \$25 is the certainty equivalent of the coin-flip.


One last test compared a variety of allocation strategies, including those seen in the popular media, for an investor who uses the shortfall-risk measure and has a target of \$1 million. The asset allocation strategies compared were:

- the dynamic programming allocation strategy;
- the allocations suggested by Malkiel (1995);
- the allocations suggested by Bogle (1994);
- the allocations suggested by Quinn (1997);
- an allocation in equities of 100% minus investor's age, with the remainder in T-bonds;
- an allocation in equities of 100% minus investor's age, with the remainder divided evenly between T-bonds and T-bills;
- 100% in equities in each year;
- 100% in T-bonds in each year; and
- 100% in T-bills in each year.

The suggested equity allocations of Malkiel, Bogle, and Quinn are described in Table 1.

For each of forty years, 500 simulated sets of annual returns were generated in the same way as those of the simulations in the previous section. Then, for each strategy, an investor who made \$5,000 annual contributions was assumed to choose asset allocations as the strategy suggested.

Figure 4
Optimal Allocations to Equities with Goal Date Thirty, Twenty, and Ten Years in the Future



This generated for each allocation strategy a distribution of 500 possible levels of wealth when the goal date was reached. It was possible to use this sample to determine a single certain amount that the investor would view as equivalent to the distribution of outcomes from that strategy, i.e., a certainty equivalent. Table 2 ranks the strategies and reports the certainty equivalents and average values for the allocation strategies considered.

The dynamic programming technique specifically links asset allocations to the individual's attitudes toward risk, and thus it provides optimal allocations for the individual. For the sample of simulated returns, these allocations produced a certainty equivalent of \$688,000 and an average wealth of \$906,000. The certainty equivalent is a benchmark against which the other plans can be compared.

The second-best strategy was to allocate 100% to equities, producing the highest average wealth at \$1,867,000. However, the distribution of outcomes was deemed sufficiently risky that the certainty equivalent was only \$659,000.

Close behind were the two most aggressive strategies advocated in the popular media, namely those recommended by Bogle and Quinn. The loss in certainty equivalents for the 100% equity allocation and the Bogle and Quinn strategies is less than 10%, quite reasonable for "off-the-rack" allocations rather than ones tailored to the specific individual's attitudes toward risk. However, under different sets of assumptions about the investor's attitudes toward risk, the rank of the various strategies would change.

The next strategies were Malkiel's strategy and two other allocation strategies often suggested in the popular media. The worst of these represents a loss of almost 20% from the optimum.

Finally, finishing last were portfolios of Treasury securities. They not only left the investor well short of his or her goal but also produced certainty equivalents less than 50% of those of the optimal allocation strategy.

Table 1
Percentage Allocated to Equities

Analyst	Goal Thirty to Forty Years Away	Goal Twenty to Thirty Years Away	Goal Ten to Twenty Years Away	Goal Less Than Ten Years Away
Malkiel	70%	60%	60%	50%
Bogle	80%	80%	70%	70%
Quinn	100%	100%	80%	80%

Table 2
A Comparison of Average Wealth and Certainty Equivalents for an Investor with a Forty-Year Savings Period

	Average Wealth in Forty Years	Certainty Equivalent
Optimal allocations	\$906,000	\$688,000
100% equities	\$1,867,000	\$659,000
Quinn allocations	\$1,484,000	\$658,000
Bogle allocations	\$1,163,000	\$641,000
Malkiel allocations	\$859,000	\$603,000
(100% minus age) in equities, remainder in T-bonds	\$790,000	\$587,000
(100% minus age) in equities, remainder split between T-bonds and T-bills	\$701,000	\$562,000
100% T-bonds	\$348,000	\$328,000
100% T-bills	\$230,000	\$229,000

Assumptions: The values in the table are for an investor making \$5,000 annual contributions, whose goal date is forty years away. The investor is assumed to have a target of \$1 million, to be moderately risk averse, and to be using the shortfall-risk method.

Conclusions

The traditional wisdom regarding asset allocation of savings toward a distant goal is that investors should be aggressive while they're young and gradually shift to more conservative investments as they age. We find that this is true for investors with specific targets in mind and that in fact, young investors should be even more aggressive than is traditionally advised. In addition, we find that optimal allocations may vary widely for investors of the same

age, depending on their targets and how much they have accumulated to date, although the "invest aggressively while you're young" rule seems to be appropriate even for the most conservative of investors. Finally, the costs of making allocation choices that are not optimal can also be significant, particularly if the investor allocates too much to "safe" investments. This suggests that great care should be taken to try to incorporate an investor's attitudes toward risk and levels of accumulated savings into the asset-allocation decision.

Endnotes

¹For a more thorough discussion of these and other problems, see Kahneman and Tversky (1979, 1984) or Thaler (1992). A more detailed introduction to expected utility functions can be found in Eeckhoudt and Gollier (1995).

²Since investors are concerned not with the number of dollars accumulated when the goal is reached, but rather with what those dollars can buy, all our analysis is in real terms (i.e., net of inflation).

³A market return of 50% is a bit extreme (in real terms it has occurred only twice since 1926), and it is used here for illustrative purposes only.

⁴The utility function used here is $U[\text{wealth}] = 10 \times (1 - e^{-\text{wealth}/400,000})$. Such an expected utility function has the property of constant absolute risk aversion, and is often used in academic models.

⁵We used the function $U[\text{wealth}] = \text{average wealth} - a(\text{semivariance of wealth})_{\text{target}}$, where $\text{semivariance} = \text{sum}(s^2)$; $s = \text{wealth} - \text{target}$ if $\text{wealth} < \text{target}$; $s = 0$ otherwise. The parameter value a was set to .001. The

semivariance method used here is consistent with the mean-lower partial moment framework proposed by Bawa and Lindenberg (1977). Surprisingly similar results (not reported here) were obtained when an exponential utility function was used.

⁶To provide some perspective, \$1 million converted to one of TIAA-CREF's variable annuities would, at today's rates, provide monthly benefits somewhere in the range of \$5,000 to \$8,000. The exact values depend on a number of conditions, including age of the investor, single-life vs. double-life, guarantee period, etc.

⁷Specifically, we used the function $U[\text{wealth}] = a(\text{semivariance of wealth})_{\text{target 1}} - b(\text{semivariance of wealth})_{\text{target 2}}$, where $a = .001$; $b = .00004$; $\text{target 1} = \$500,000$; and $\text{target 2} = \$1 \text{ million}$.

Bibliography

- Bawa, V., and E. Lindenberg. "Capital Market Equilibrium in a Mean, Lower Partial Moment Framework." *Journal of Financial Economics* 5, no. 2 (November 1977): 189-200.
- Bogle, J. *Bogle on Mutual Funds*. Homewood, Illinois: Richard D. Irwin, 1994.
- Eeckhoudt, L., and C. Gollier. *Risk: Evaluation, Management, and Sharing*. Trans. V. Lambson. London: Harvester Wheatsheaf, 1995.
- Ibbotson Associates. *Stocks, Bonds, and Inflation 1996 Yearbook*. Chicago: Ibbotson Associates, Inc., 1996.
- Kahneman, D., and A. Tversky. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47, no. 2 (March 1979): 263-91.
- . "Choices, Values, and Frames." *American Psychologist* 39 (1984): 341-50.
- Landsburg, S. *The Armchair Economist*. New York: Simon and Schuster, 1994.
- Malkiel, B.G. *A Random Walk down Wall Street*. 6th ed. New York: W. W. Norton & Co., 1995.
- Quinn, J. *Making the Most of Your Money*. New York: Simon and Schuster, 1997.
- Thaler, R. *The Winner's Curse*. Princeton, New Jersey: Princeton University Press, 1992.

