

Why Should Older People Invest Less in Stocks Than Younger People?

Ravi Jagannathan
Visitor
Research Department
Federal Reserve Bank of Minneapolis
and Piper Jaffray Professor of Finance
Carlson School of Management
University of Minnesota

Narayana R. Kocherlakota
Economist
Research Department
Federal Reserve Bank of Minneapolis

Abstract

Financial planners typically advise people to shift investments away from stocks and toward bonds as they age. The planners commonly justify this advice in three ways. They argue that stocks are less risky over a young person's long investment horizon, that stocks are often necessary for young people to meet large financial obligations (like college tuition for their children), and that younger people have more years of labor income ahead with which to recover from the potential losses associated with stock ownership. This article uses economic reasoning to evaluate these three different justifications. It finds that the first two arguments do not make economic sense. The last argument is valid—but only for people with labor income that is relatively uncorrelated with stock returns. If a person's labor income is highly correlated with stock returns, then that investor is better off shifting investments toward stocks over time.

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Most financial planners advise their clients to shift their investments away from stocks and toward bonds as they age. For example, in *The Wall Street Journal Guide to Planning Your Financial Future*, Kenneth Morris, Alan Siegel, and Virginia Morris (1995, p. 7) tell people to make sure that the percentage of wealth they have in bonds is no more than their age. Similarly, Jane Bryant Quinn (1991, p. 489), investment columnist for *Newsweek*, tells investors to “tip toward higher risks if you . . . are young.” And in the classic book *A Random Walk Down Wall Street*, Burton Malkiel (1996, p. 411) advises “more common stocks for individuals early in the life cycle and more bonds for those nearer to retirement”; he says that “the longer the time period over which you can hold on to your investments, the greater should be the share of common stocks in your portfolio” (Malkiel 1996, pp. 404–405).

Despite their general agreement that investors should switch from stocks to bonds as they age, financial planners give different reasons for recommending this investment policy. At least three reasons are commonly offered. First, many financial planners argue, as does Malkiel (1996, p. 403), that “a substantial amount . . . of the risk of common-stock investment can be eliminated by adopting a program of *long-term* ownership,” and, of course, older people don’t have as many years ahead of them as do younger people. Second, some financial planners emphasize that asset allocation is often shaped by the necessity of meeting relatively large obligations in midlife, such as college tuition for children. To meet these financial targets, investing a lot in stocks may be necessary for a while, but not after enough resources have accumulated. And finally, some financial planners point out, as again Malkiel (1996, p. 400) does, that a younger person “can use wages to cover any losses from increased risk” while an older person cannot.

In this article, we use standard economic models of investor behavior to evaluate each of these explanations.¹ We conclude that the low long-term risk of stocks explanation and the targeting explanation have little validity. The only explanation that holds up as solid justification for the stock holding advice is the fact that younger people have many years of wages available to them while older people do not.

We begin by documenting that, as Malkiel and others state, stocks are much more likely to outperform bonds over long horizons than over short horizons; in this sense, stocks become less risky over longer horizons. However, we show that this fact is irrelevant for investors, for two reasons. One reason is obvious: if investors can rebalance their portfolios over time, a long horizon is basically the same as a short horizon; what matters for investment decisions is the length of time between rebalancing, not the investment horizon itself. The other reason for the irrelevance of low long-term risk is subtler. Even if investors can’t rebalance their portfolios, they have to be concerned about the potential for enormous losses that can be incurred by holding stocks over long periods of time. For example, over a 30-year period, the events of 1929 can occur 30 times; those same events can only occur once in a one-year period. While having 30 such poor years in a row may be exceedingly unlikely, we show that according to standard economic models of investor choice, investors are concerned about the magnitude of these potential losses,

not just their probability. Standard models predict that because of this concern, investors will split their wealth between stocks and bonds in the same way, independent of the length of their investment horizon. We conclude that the reduction in the riskiness of stocks over longer horizons does not justify the common advice of financial planners.

We look next at the explanation that asset allocation is often shaped by large needs in midlife—some targets that must be hit, such as enough financial wealth to pay for college tuition for children. We find that when confronted with such a need, some investors will indeed find their best move is to switch from stocks to bonds over time. Generally, though, such a switch is extremely dramatic, not the gradual reductions typically recommended by financial planners. Moreover, whether investors actually switch toward bonds or away from bonds as they age depends crucially on the size of their target, their initial wealth, and the loss associated with failing to hit the target. Since an optimal plan is so dependent on investor-specific variables, we conclude that this explanation does not justify financial planners generally recommending risk reduction as investors age.

Finally, we consider the explanation that the life-cycle behavior of labor income shapes investor behavior. We find that there is a good economic justification for this explanation. When investors are young, they have a long stream of future income. As they age, this stream shortens, so the value of their human capital falls. (If labor income is rising over time, the value of human capital may rise initially, but eventually it has to fall because the amount of time left before retirement starts to shrink at a very fast rate.) The best way for investors to respond to this situation is to shift the risk composition of their financial wealth in order to offset the decline in the value of their human capital. For most people, labor income either is risk free or is dominated by person-specific risk that is only weakly correlated with stock returns. So most investors need to shift their financial wealth toward bonds and away from stocks as they age in order to make up for the loss in human capital. We conclude that substituting for lost labor income is the only valid reason for financial planners’ advice that clients shift their portfolios toward relatively riskless instruments as they age.

The mathematics behind our analysis is hardly new; it was first derived in Robert Merton’s (1971) classic paper.² Why do we find it necessary to reemphasize the lessons of his work? We have a very practical reason: today many more investors than ever before are able to control their own asset allocations. This can only be done intelligently if one knows the basis of the financial planners’ advice. For example, suppose a young investor has an income stream that is highly correlated with stock returns. Financial planners generally would advise this person to invest less in stocks as time passes. We show that this investor should not do that, but rather should invest more in stocks in order to make up for the loss of labor income.³

Risk in the Long Run

First we consider the argument that younger investors should invest more in stocks than older investors because stocks are less risky over longer investment horizons. We show that there is certainly a sense in which stocks are less risky over longer horizons. However, we also show

that this does not mean that investors will be better off if they invest significantly more in stocks when their investment horizon is longer.

We begin by documenting the historical behavior of returns on stocks and U.S. Treasury bills (T-bills) over the period 1926–90 (as reported in Ibbotson Associates 1992). During these 65 years, the average annual real return to the stocks of the 500 large firms in Standard & Poor's stock price index (the S&P 500) was about 8.8 percent per year. (The average of the logarithm of the gross real return was 6.5 percent.) Over the same period, T-bill real returns averaged about 0.6 percent. Thus, stocks earned a remarkable 8.2 percentage point annual premium over T-bills. Stocks, of course, were much more variable: the standard deviation of the annual real return to the S&P 500 was about 21 percent. (The standard deviation of the logarithm of the gross real return was 20 percent.) In contrast, the standard deviation of the annual real return to T-bills was only about 4.4 percent.

Following Malkiel (1996), investment advisers generally emphasize two features of these data. First, bills outperformed stocks in 20 years out of a possible 65. Second, the sample has 46 possible blocks of 20 consecutive years. In none of these blocks did bills outperform stocks. These facts are generally interpreted as saying that while stocks are risky over short horizons, they are guaranteed to outperform bills over a 20-year period.

Unfortunately, this conclusion is somewhat premature. While the sample has 46 possible blocks of 20 consecutive years, it has only 3 nonoverlapping (independent) blocks of 20 years. This means that the sample itself contains little direct information about the long-run performance of stocks compared to bills. We need to augment the sample information with information from economic theory and construct a statistical model of stock and bond returns. We can then use that model to address questions about the long run.

We obtain this additional theoretical information from what is known as the *random walk hypothesis*. This theory is based on the following simple logic. Stock prices reflect all available information, which means that stock prices change only if news arrives. News is by definition unpredictable. Hence, to a first-order approximation, stock price changes are unpredictable.

We embed this theoretical reasoning in a statistical model by assuming that stock returns are independent and identically distributed over time.⁴ We then assume that logged stock returns are normally distributed,⁵ with mean, or *expectation*, and standard deviation (μ_S and σ_S , respectively) equal to their sample values, 6.5 percent and 20 percent. We ignore the relatively small variability of T-bill returns and assume that, within the model, bond returns are constant at 0.5 percent per year. As is common in modern dynamic economics, we model households as having *rational expectations*: they know that stock and bond returns behave in the way described by our statistical model.

Using this statistical model, we can assess the claim that stocks are guaranteed to outperform bonds over long enough horizons. Suppose someone has a dollar to invest. If he or she were to put all of it into stocks, then the logarithm of the amount of wealth this investor would have after T years, $\ln(W_T)$, would be random, with mean $\mu_S T$ and standard deviation $\sigma_S T^{1/2}$. But if the investor were to put

the dollar into T-bills, then $\ln(W_T)$ would be nonrandom and equal to $\ln(1.005)T$.

Notice that the mean of the difference between the two portfolios' payoffs increases linearly with T . But the standard deviation of the stock portfolio's payoff increases much more slowly—only linearly with $T^{1/2}$. Thus, when T increases from 1 to 30, the mean difference increases by a factor of 30, while the standard deviation increases by a factor of only 5.5. This means that for large T , the mean of the difference between the two portfolios' payoffs is going to be large and positive relative to the standard deviation of this difference. Hence, for large T , the difference in the payoffs is very unlikely to be negative.

This intuition is illustrated quantitatively in Chart 1. It shows that according to our statistical model, over a one-year period, the stock portfolio outperforms the bond portfolio with a probability of approximately 0.6. However, the probability of getting a better return with stocks over a 30-year period is 0.95. Thus, our statistical model does imply that over long periods of time, bonds are highly unlikely to outperform stocks;⁶ yet the model also implies that there is some (albeit small) probability that bonds will outperform stocks even over 30-year horizons.

We did a test to check the ability of our statistical model to fit the long-run properties of the data. First, we simulated 1,000 samples of length 65. In each of these samples, we looked at the 46 possible blocks of 20 consecutive years. In 500 of the samples, bonds failed to outperform stocks in any of the 46 possible 20-year periods.

Data like those displayed in Chart 1 are often used to justify the advice that younger people should invest more in stocks than older people: because stocks are more likely to do better than bonds over the long haul, financial advisers recommend investing more in stocks when the investment horizon is long. But this reasoning ignores two crucial aspects of optimal portfolio allocation. First, investors can readjust their portfolios over time. With an ability to rebalance, how is a long horizon different from a short horizon? Second, most households are concerned not just with the probability of loss, but also with the magnitude of the loss. Now we evaluate the relevance of Chart 1, given these two issues.

To understand the importance of the first issue, consider the decision problem of a household which has $\$W_0$ available today to invest. The household's goal is to maximize the expected value of a utility function $U(W_T)$, where W_T is the amount of wealth the household will accumulate over T periods.⁷ The form of the utility function U is an important determinant of the household's behavior. Standard dynamic economic models assume that households have objective functions⁸ with constant relative risk aversion, so that

$$(1) \quad U(W) = W^{1-\gamma}/(1-\gamma)$$

where the parameter γ represents the level of risk aversion. According to this objective function, households with $\$10,000$ to invest will split their wealth (W) between stocks and bonds in the same way as households with $\$100,000$ to invest. When the parameter γ is high, households are more risk averse and will invest less money in stocks. Generally, economists restrict the parameter γ to lie

between 0 and 10. (For a closer look at risk aversion, see the accompanying box.)

We assume that the household can invest in two different accounts. One is a stock mutual fund with annual real returns r_t^s . As before, we assume that returns are independent and identically distributed over time and that logged returns, $\ln(1+r_t^s)$, are normal with mean 6.5 percent and standard deviation 20 percent. The other account is a bond mutual fund that pays a constant real return $r^b = 0.5$ percent. The household chooses its stock holdings s_t and its bond holdings b_t in each period so that it solves this maximization problem:

$$(2) \quad \max_{(S_t, B_t)_{t=0}^{T-1}} E\{(W_T)^{1-\gamma}\}/(1-\gamma)$$

subject to

$$(3) \quad W_T = S_{T-1}(1+r_T^s) + B_{T-1}(1+r^b)$$

$$(4) \quad S_t + B_t \leq (1+r_t^s)S_{t-1} + (1+r^b)B_{t-1}$$

for $1 \leq t \leq T-1$ and subject to

$$(5) \quad S_0 + B_0 \leq W_0$$

for W_0 given. In this problem, S_t and B_t are the amounts of money that the household invests in stocks and bonds, respectively, in period t . The household begins life with W_0 .

We can solve the household's problem backward. At the end of period $T-1$, the household has W_{T-1} dollars. It chooses the share $s_{T-1} \equiv S_{T-1}/W_{T-1}$ of this wealth to invest in stocks so that it solves this maximization problem:

$$(6) \quad \max_{s_{T-1}} E\{[s_{T-1}W_{T-1}(1+r_T^s) + (1-s_{T-1})W_{T-1}(1+r^b)]^{1-\gamma}\}/(1-\gamma).$$

Here, the expectation averages the household's payoff from its portfolio over the possible realizations of r_T^s . The key feature of this optimization problem is that the household's choice of s_{T-1} is independent of its wealth W_{T-1} ; the household's utility function has constant relative risk aversion. In fact, because r_t^s is independent over time, the household's choice of s_{T-1} equals s^* in all states, where s^* satisfies this first-order condition:

$$(7) \quad E\{[s^*(1+r_t^s) + (1-s^*)(1+r^b)]^{-\gamma}(r_t^s - r^b)\} = 0.$$

This first-order condition equates the marginal benefits of investing in stocks and bonds.

Now that we know how the household will solve its problem in period $T-1$, we can work backward to figure out its optimal portfolio in period $T-2$. In that period, the household realizes that in period $T-1$, it will invest a share s^* of its wealth in stocks. It takes this into account when deciding how much to invest in period $T-2$ and therefore solves this problem:⁹

$$(8) \quad \max_{s_{T-2}} E\{[s_{T-2}(1+r_{T-1}^s) + (1-s_{T-2})(1+r^b)] \times [s^*(1+r_T^s) + (1-s^*)(1+r^b)]\}^{1-\gamma}/(1-\gamma).$$

Because r_{t-1}^s is independent over time, we can separate the expectation of the product into the product of the expectations. That is, we can rewrite the household's problem as

$$(9) \quad [s^*(1+r_t^s) + (1-s^*)(1+r^b)]^{1-\gamma} \times \max_{s_{T-2}} E\{[s_{T-2}(1+r_{T-1}^s) + (1-s_{T-2})(1+r^b)]\}^{1-\gamma}/(1-\gamma).$$

The household's objective is thus the same in period $T-2$ as in period $T-1$. So it is optimal for the household to set s_{T-2} equal to s^* . We can keep rolling this logic backward in time to conclude that it is optimal for s_t to equal s^* in every date and state.

Intuitively, we know that the investment decision of a household with constant relative risk aversion is always independent of the household's current wealth level—in fact, that's essentially the definition of *constant relative risk aversion*. Moreover, we assume that stock returns are independent and identically distributed over time, so that the household's current beliefs about next period's returns are always the same. Thus, long investment horizons are no different from short investment horizons as long as households are making decisions at regular intervals.

In the real world, transaction costs lead households to change their portfolios only infrequently. Let's see how this restriction on household behavior affects portfolio allocation. Suppose that in our model, households can only split their initial wealth W_0 between stocks and bonds; after that investment decision, they can never move resources from one fund to another. Then, after T periods, the household's wealth is random:

$$(10) \quad W_T = [S_0(1+r_1^s)(1+r_2^s)(1+r_3^s) \dots (1+r_T^s)] + (W_0 - S_0)(1+r^b)^T$$

where S_0 is the amount of wealth that the household puts into the stock market account today. Again, the optimal S_0/W_0 is independent of W_0 . The household chooses S_0 so as to maximize $E\{(W_T)^{1-\gamma}\}/(1-\gamma)$. This maximization problem is easy to solve using numerical methods.¹⁰

Table 1 describes how the share of wealth that a household will invest in the stock market s depends on the horizon T and on the risk aversion parameter γ . We find that when households have coefficients of relative risk aversion at least as high as 2, then the portion s is virtually independent¹¹ of T . At first, this result seems paradoxical; after all, as T grows, so does the probability that stocks will outperform bonds. But there is another effect. Suppose each time a coin is flipped, an individual receives \$2 if the coin shows heads, but loses \$1 if the coin shows tails. If the coin is flipped once, the individual can lose up to \$1. However, if the coin is flipped 30 times, the individual can lose up to \$30 (although the probability of doing so is quite small). Table 1 says that for households with constant relative risk aversion, the increased potential for very poor performance associated with long-term stock investments almost exactly offsets the increased potential for very good performance.

In fact, in a well-defined sense, portfolios composed entirely of stocks may become less attractive over longer horizons to households that are sufficiently risk averse. Suppose households are asked how much they would have to invest in a stock portfolio to be indifferent between it and \$1 invested in the household's optimal portfolio (assuming no rebalancing) or in a bond portfolio. The answer to this question is in Table 2. It shows that over a 40-year horizon, a household with coefficient of relative risk aversion

equal to 5 prefers \$1 invested in the optimal portfolio to anything less than \$3.86 invested in the stock portfolio. Even more surprising, the household prefers \$1 invested in a bond portfolio to anything less than \$2.22 invested in the stock portfolio—even though the stock portfolio outperforms the bond portfolio with a probability of 97 percent. Risk averse households are clearly highly concerned with the possibility for extreme losses when stocks don't do well, even though this is a very unlikely event.

Thus, despite the apparent attractiveness of stocks over the long haul, the length of the investment horizon does not greatly affect the investment decisions of households with constant relative risk aversion utility functions. If these investors can freely rebalance their portfolios, they are essentially facing a sequence of one-period decision problems, regardless of the length of their investment horizons. Even if they can't rebalance freely, they regard the increased downside risk generated by longer investment horizons as a counter to the increased upside potential depicted in Chart 1.

Targeting

Now we consider the argument that investors generally should reduce their stock holdings over time because they are saving for a particular target level of wealth (especially for midlife obligations like college tuition for their children).

One way to capture this idea is to assume that the household derives utility from final wealth W_T according to a function that is defined as

$$(11) \quad U(W_T) = \begin{cases} (W_T - \bar{W})^{1-\gamma}/(1-\gamma), & \text{if } W_T \geq \bar{W} \\ -L & , \text{if } W_T < \bar{W} \end{cases}$$

According to this objective function, the household treats \bar{W} as a target. It receives utility from any amount exceeding \bar{W} according to the constant relative risk aversion function described earlier. Failing to achieve \bar{W} results in a negative amount of utility. For now, we set this loss $L = \infty$, so that no potential gain can outweigh even the slightest probability of failing to reach the target. While seemingly extreme, this specification of L is common in economics because it is the only value of L for which the function in (11) is concave.

We have seen that if a household can rebalance its portfolio costlessly and it has no target level of wealth, it will always keep the same proportion of wealth in stocks. How does targeting affect this result? To answer this, suppose first that $\bar{W}/(1+r^b)^T \geq W_0$. This condition means that the household can exceed its target by simply investing all of its funds in the risk-free asset. The household then follows a two-step investment procedure. First, it always invests enough money in bonds so as to definitely achieve the target (because not achieving it has such dire consequences). Second, the household invests a constant share s^* of any additional money into stocks [where s^* satisfies the first-order condition (7)].

This immediately means that the household invests a smaller share of its wealth in stocks than it would have if it did not have a target. But as time goes by, the household accumulates wealth, and because it invests in both stocks and bonds, its wealth typically grows faster than the inter-

est rate. Hence, as the household approaches retirement, it has to use proportionately fewer of its resources to be sure of achieving its target \bar{W} . This frees up more funds for stock investment. It follows that over time a targeting household tends to increase the share of its wealth in stocks, not decrease it as financial planners advise.

All of this analysis presumes that the target can be reached by using only bonds. Suppose instead that the household is faced with a target that cannot be reached in that way; that is, its target $\bar{W} > W_0(1+r^b)^T$. This means that no matter what investment strategy the household uses, there is always some probability of not achieving the target and so getting utility equal to $-\infty$. Therefore, when the target is so large relative to its initial wealth, the household views all strategies as equally bad and is indifferent among all possible investment strategies.

This implication may strike many readers as somewhat strange. In response to that likely reaction, we explore in the Appendix what happens if we change the loss L to some less extreme amount. There we show that situations exist in which, if L equals zero, then the household tends to invest a smaller share of its resources in stocks over time. However, this behavior is only optimal for a narrow range of target-to-wealth ratios. Because of this lack of robustness, we conclude that targeting does not justify the advice financial planners generally give their clients, to lower the share of their portfolios in stocks over time.

Labor Income

So far, we have implicitly modeled the household as receiving all of its income in the form of interest, capital gains, and dividends. This assumption probably accurately describes only a few actual households in the United States: most derive much of their income from working in the labor force. Now we broaden our model and discover that this other type of income provides the most convincing justification for the stock holding advice of financial planners.

Suppose that a household receives a salary y_t in period t and begins its working life with an initial amount of stocks and bonds. The household can invest its salary in both stocks and bonds and can costlessly rebalance its portfolio at the end of each period. Here, as before, we assume that the household's objective is to maximize the expectation of the utility function with constant relative risk aversion, $(W_T)^{1-\gamma}/(1-\gamma)$, where W_T is its wealth at the end of period T . We can now write the household's problem as

$$(12) \quad \max_{\{S_t, B_t\}_{t=0}^T} E\{(W_T)^{1-\gamma}\}/(1-\gamma)$$

subject to

$$(13) \quad W_T = y_T + S_T(1+r_T^s) + B_{T-1}(1+r^b)$$

$$(14) \quad S_t + B_t \leq y_t + (1+r_t^s)S_{t-1} + (1+r^b)B_{t-1}$$

for $1 \leq t \leq T-1$ and subject to

$$(15) \quad S_0 + B_0 \leq W_0$$

for W_0 given. In this problem, as before, S_t and B_t are the amounts of money that the household invests in stocks and bonds, respectively, in period t .

The solution to this problem depends on the fact that the household has two types of wealth. One is its *financial*

wealth (FW), the money that it has invested in stocks and bonds. Mathematically, we can express this as

$$(16) \quad FW_t \equiv y_t + (1+r_t^s)S_{t-1} + (1+r_t^b)B_{t-1}.$$

The other type of wealth is the household's *labor wealth*, its sequence of future salary payments, $(y_{t+1}, y_{t+2}, y_{t+3}, \dots)$. The risk characteristics of this second type of wealth play a crucial role in determining the household's optimal split of its financial wealth between stocks and bonds. Suppose, for example, that the salary y_t equals a constant y in every period. Then the household's salary payments are, in risk terms, equivalent to the payments it would receive from a large risk-free annuity. Since the household is already holding a large amount of wealth in a risk-free asset, it will compensate by investing more of its resources in stocks.¹²

The mathematical formalization of this basic intuition is useful and simple. If the household's salary is constant over time, then at the end of any period t , the present value of its future salary payments equals

$$(17) \quad PVY_t = [y/(1+r^b)] + [y/(1+r^b)^2] + \dots + [y/(1+r^b)^{T-t}].$$

(Note that we can discount the future payments by the risk-free rate because they are risk free.) Earlier, in equation (7), we saw that a household with constant relative risk aversion wants to have a constant share s^* of its total wealth in stocks, where s^* satisfies the first-order condition

$$(18) \quad E\{[s^*(1+r_t^s) + (1-s^*)(1+r_t^b)]^{-\gamma}(r_t^s - r_t^b)\} = 0.$$

Now, with two types of wealth, at the end of any period t , the household will set S_t to equal

$$(19) \quad S_t = s^*(\text{Total Wealth}) = s^*(FW_t + PVY_t).$$

The crucial aspect of this last formula is that the ratio of financial wealth to labor wealth is not constant over time. Consequently, even though the share of total wealth held in stocks does not change over time, the share of financial wealth held in stocks does:

$$(20) \quad S_t/FW_t = s^*[1 + (PVY_t/FW_t)].$$

Since FW_t is random, we cannot state with certainty how S_t/FW_t changes over time. However, we can infer how S_t/FW_t is likely to change. Two forces are at work. One is that, as the household gets closer to retirement, it has fewer salary payments left to receive, and the present value of its future salary payments falls. The other force at work is that, as time passes, the household accumulates more stocks and bonds, so FW_t tends to rise. Because of these two forces, the household's optimal plan will typically be to reduce S_t/FW_t over time.¹³

This analysis is purely qualitative. Chart 2 offers a feel for the magnitudes of the changes in stock holdings involved. It depicts how the share of financial wealth held in stocks changes over time for a household that has \$20,000 per year available for investment from its salary income and that has a risk aversion coefficient equal to 5. (Returns from stocks and T-bills follow the statistical model we discussed earlier.) The household begins its working life

at age 25 with no financial wealth; the chart depicts what happens to its path of stock holdings from age 36 until its last investment decision before retirement at age 64. Of course, for any particular household, the path of stock holdings is random (because stock returns fluctuate over time). Chart 2, therefore, displays the median path of stock holdings over 1,000 randomly drawn time paths of stock returns. It shows that at age 36, the household should have nearly 85 percent of its financial wealth in stocks; by age 64, this share drops below 40 percent.¹⁴

While striking, this analysis has an obvious weakness: households do not actually receive constant salaries. Instead, from year to year, a typical household does not know what the real value of its income will be. Random salaries mean that the composition of a household's portfolio over time depends crucially on the degree of comovement between the household's salary and the return to the stock market.

To make this dependence clear, consider the extreme example of the manager of a mutual fund. Suppose that the growth rate of the manager's income always equals the growth rate of aggregate dividends. Then, from a risk point of view, the manager's future labor income is essentially equivalent to having a lot of money invested in stocks. As the manager approaches retirement, the value of this future labor income falls. That implicitly reduces the share of the manager's wealth that is highly correlated with the stock market. The manager uses financial assets to compensate for this change by investing more in stocks over time.

This story suggests that households should invest more in stocks over time. Of course, the premise of the reasoning is that the growth rate of household salaries is highly correlated with dividend growth, which is unrealistic for many people; after all, not everyone is a mutual fund manager. In fact, based on their examination of microeconomic evidence on household labor income, John Heaton and Deborah Lucas (1996) find that most households have incomes that are not highly correlated with the performance of the stock market.¹⁵ With this statistical characterization, even though wages are risky, households think of bonds as a closer substitute than stocks for their labor income.¹⁶ Consequently, the optimal plan for many households is the advice financial planners give: when you're young, offset current and future labor income by holding a lot of stocks; as time passes and fewer periods remain in which to earn labor income, compensate by increasing bond holdings.¹⁷

Conclusion

The life-cycle advice of financial planners is now so widely known that it can be termed *folk wisdom*: older people should invest less in stocks than younger people do. But why should they? We have here used standard economic reasoning to evaluate three reasons that are offered by many financial planners, and we have shown that only one of them makes economic sense.

That one has to do with the fact that as investors age, they have fewer years of labor income ahead of them. If—as is true for most people—an investor's labor income is not directly correlated with stock returns, then our economic analysis concludes that the investor should follow the financial planners' advice: as time passes, shift more financial wealth out of stocks and into bonds.

That's not good advice for investors who are not like most people, though. If an investor's labor income is high-

ly correlated with stock returns, then our economic analysis demonstrates that the investor is likely to be better off ignoring the common advice of financial planners. Instead, such an investor should do the opposite of what financial planners say: as time passes, shift more financial wealth out of bonds and into stocks.

*The authors thank Karen Hovermale for valuable research assistance; Rao Aiyagari, Lee Ohanian, Víctor Ríos-Rull, and Kathy Rolfe for their comments; and John Heaton for helpful conversations. Kocherlakota thanks John Kennan, Barbara McCutcheon, Sergio Rebelo, and Chuck Whiteman for many discussions in the distant past about the issues in this article.

¹Other justifications for this type of variation in stock holdings over the life cycle can, of course, be constructed. But we choose to focus on those most often used by financial planners.

²The mathematical analysis is also restated as a special case of the analysis in Zvi Bodie, Robert Merton, and William Samuelson's (1992) paper. They are more explicit than Merton (1971) in discussing the implications of his original analysis for portfolio dynamics over the life cycle. Our conclusions about the role of labor income essentially mirror theirs.

³We should note that our model consistently overpredicts the amount of stock holdings of households at every stage of the life cycle. This is because we abstract from taxes, real estate investment, short-sale constraints, borrowing restrictions, and endogenous labor supply (among other things). We know that ignoring these elements affects the predictions of our model for the quantitative path of stock holdings over the life cycle. But our goal here is limited: we simply want to determine qualitatively whether any of the explanations for the common investment advice is robust.

So far, no economic model has satisfactorily explained the low level of stock holdings given the large difference in average returns between stocks and U.S. Treasury bills (T-bills). This is essentially a partial equilibrium manifestation of the equity premium puzzle: no satisfactory general equilibrium model is simultaneously consistent with the low variability of per capita consumption growth and the wide spread between average stock and T-bill returns. See the article by Narayana Kocherlakota (1996).

⁴The theoretical argument actually implies only that investors have no information available that allows them to forecast mean returns. We strengthen this assumption to independence. Eugene Fama and Kenneth French (1988) present evidence that stock returns have a predictable component. However, the sampling errors associated with estimates of predictability are very large. [See the work of Robert Hodrick (1992, Table 4, Panel D).] Given the theoretical argument and the lack of empirical evidence, a conservative view for planning purposes is to assume no predictability.

⁵The assumption of normality is not a bad approximation for the empirical distribution of stock returns (except for some events in the left tail of the distribution). More important, the normality assumption is made purely for analytical convenience; using the empirical distribution instead would not affect any of our conclusions.

⁶Actually, this is true in any model in which stocks have a higher population mean return than bonds.

⁷Throughout this article, we ignore the consumption/saving decision of the household and focus only on the portfolio allocation problem. Merton (1971) shows that the household's portfolio allocation decision can be found separately from its consumption/saving decision if returns are independent and identically distributed over time, if the household has constant relative risk aversion (or what we call *concave targeting*) preferences which are separable from leisure, and if the household faces no short-sale constraints or uninsurable income risk.

⁸This assumption is made for one key reason: even though per capita consumption has gone up by about eight times in the United States in the last 130 years, real rates of return have remained relatively steady. This would not be true if household objective functions exhibited increasing or decreasing relative risk aversion.

⁹Note that we do not constrain s_{T-1} in any way. Except for the nonconcave objective discussed in the Appendix, we do not impose short-sale or borrowing constraints. We conjecture that imposing these would not change the flavor of our results; but actually solving the optimization problems while imposing these constraints is a problem beyond our scope here.

¹⁰Specifically, we solve the problem by approximating the standard normal distribution by a 10-point discrete distribution that has the same first 19 moments as the standard normal. Using an approximating 9-point discrete distribution delivers much the same results.

¹¹Note that in Table 1, households that have low risk aversion (a coefficient less than 4) prefer to invest more in stocks over longer horizons, while those with high risk aversion (a coefficient greater than or equal to 4) prefer to invest less in stocks over longer horizons. This reversal in behavior happens because the highly risk averse households are more concerned with the possibility for more dramatic losses over longer horizons.

¹²Throughout this section, we assume that households cannot adjust their earnings in response to stock market performance. Bodie, Merton, and Samuelson (1992) show that if households can freely choose effort at every point in time, then they choose to increase effort whenever stocks do poorly. This induces an endogenous negative correlation between stock returns and labor earnings; correspondingly, households tend to invest relatively more in stocks and to decrease the portion invested in stocks more rapidly when they are younger than when they are older.

José Víctor Ríos-Rull (1994) considers a model in which households have nonseparable preferences over consumption and leisure and labor earnings are perfectly positively correlated with stock returns. His model predicts that older households hold more stocks than younger households.

¹³Here is a precisely worded theorem which summarizes that analysis. Suppose that $y_t = y$ for all t and that $\text{prob}(r_t^s > 0) > 1/2$ (as is true in our statistical model). Then $\text{prob}(s_t/FW_t < s_{t-1}/FW_{t-1}) > 1/2$. More generally, suppose that the sequence of salary payments is risk free, with a constant growth rate g . It is easy to show that there exists \hat{t} such that $PVY_t/PVY_{t-1} < 1$ for $t > \hat{t}$. Then the following statements are both true:

- If $\text{prob}(r_t^s > 0) > 1/2$, then $\text{prob}(s_t/FW_t < s_{t-1}/FW_{t-1}) > 1/2$ for $t > \hat{t}$.
- If $\text{prob}(r_t^s > g) > 1/2$, then $\text{prob}(s_t/FW_t < s_{t-1}/FW_{t-1}) > 1/2$ for all t .

¹⁴Actually, Chart 2 dramatically understates the size of the decline in stock holdings. Early in a household's working life, the optimal financial plan for the household is to borrow heavily and invest the proceeds in stocks. Indeed, at age 26, the median household's optimal plan is to invest about \$340,000 in stocks—having borrowed \$320,000 of that amount. (Of course, that \$320,000 is but a small portion of the present value of the household's labor income.) While seemingly extreme, this is certainly in keeping with the standard financial planners' advice to invest as much as possible in stocks when young. (Solving the problem while imposing borrowing constraints is much more challenging and is certainly beyond the scope of this article. However, we doubt that imposing such constraints would change the shape of Chart 2.)

¹⁵More specifically, Heaton and Lucas (1996) look at a sample of households from the Panel Study in Income Dynamics. For each household, they use time series data to get a sample estimate of the correlation of labor income growth with the return to the Center for Research in Security Prices value-weighted portfolio. Heaton and Lucas find that in this set of household-specific correlation estimates, the median is 0.02.

Ravi Jagannathan and Zhenyu Wang (1996) argue that portfolios of stocks can be constructed that have returns which are positively correlated with aggregate labor income growth (even conditional on aggregate market conditions). Note that these findings are not inconsistent with those of Heaton and Lucas (1996). The correlation for a typical household may be much lower than the correlation of aggregate or per capita labor income growth with the value-weighted return. Individual household labor income features a large amount of job-specific risk (disability and bad matches, for example) that is essentially uncorrelated with the stock market; this job-specific risk cancels when individual labor income is added up across individuals to create aggregate labor income. Of course, for our purposes, household or individual labor income is the variable of interest.

¹⁶In the Appendix, we prove that if the equity premium is large enough and the precautionary motive is small enough or both, then in a one-period setting, investors view risk-free bonds as substitutes for risky labor income that is independent of stock returns.

¹⁷In the Appendix, we describe some numerical simulations in a three-period environment that confirm this intuition. Neither Merton's (1971) analysis nor that of Bodie, Merton, and Samuelson (1992) treats this case in which labor income risk is not perfectly correlated with stock market risk. In the current economic literature, we know of no numerical or analytical solutions in this incomplete markets case for households that live for a large but finite number of periods.

Appendix Numerical and Analytical Details

Here we provide more detailed arguments for some points made in the preceding paper.

The Targeting Mentality

First we demonstrate that some targeting situations exist in which stock investments are likely to shrink over time.

We examine the behavior of a household that is two years from retirement and faces a nonconcave objective function. The household solves the following problem:

$$(A1) \quad \max EU(W_2)$$

subject to

$$(A2) \quad W_2 \leq S_1(1+r_2^s) + B_1(1+r^b)$$

$$(A3) \quad S_1 + B_1 \leq S_0(1+r_1^s) + B_1(1+r^b)$$

$$(A4) \quad S_0 + B_0 \leq W_0$$

for W_0 given and, for $t = 0, 1$, subject to

$$(A5) \quad S_t \geq 0$$

$$(A6) \quad B_t \geq 0.$$

Note that in this problem, we impose short-sale and borrowing constraints to ensure the existence of a maximum. We assume that the stock return r_t^s has three equally likely realizations: 0.20, 0.01, and -0.18 ; the bond return r^b is equal to 0.005. (The low average return of the stock is necessary in order to generate any nondegenerate dynamics in portfolio holdings because in this

exercise we assume relatively low levels of risk aversion.) The crucial aspect of the problem is that $U(W_2)$ equals $(W_2 - \bar{W})^{1/2}$ if $W_2 \geq \bar{W}$ and 0 otherwise. This function is, of course, nonconcave.

We solve the one-period version of this problem by solving for the unique S_1 that satisfies the first-order conditions and then comparing the utility of that point with the utility of investing all funds in stocks. (Investing all funds in bonds is clearly suboptimal.) Given this one-period solution (which can be made very precise), we solve the two-period problem using a grid search.

The accompanying chart depicts the probability that $S_1/W_1 < S_0/W_0$, for various values of the target/wealth ratio \bar{W}/W_0 . For small values of that ratio, the household follows the behavior described in the preceding paper: the household invests enough of its resources in bonds to guarantee achieving the target and splits its remaining funds between stocks and bonds. We have seen that this strategy implies that if the stock return is higher than the bond return in any period, then the share of resources invested in stocks rises. This happens with probability 2/3 for the specification of returns in this example.

For larger values of the target/wealth ratio, however, the household invests a lot of resources in stocks in the first period in order to maximize the probability of achieving its target. If stock returns are sufficiently high that it gets over the hump after the first period, then it invests relatively little in stocks in the second period in order to avoid getting a low return and sliding below the target. If the target is somewhat low, then the household has a good chance (2/3) of beating the target after one period. If the target/wealth ratio is high, then the household beats the target with only probability 1/3. The chart shows that the household reduces the portion of its wealth in stocks with a high probability only over a small range of ratio values.

Risk-Free Bonds as Substitutes for Labor Income

Analytically

Now we prove analytically that under certain circumstances investors view risk-free bonds as substitutes for risky labor income that is independent of stock returns.

Let y be random labor income, and let r^s be random stock returns. Assume that the two random variables are independent and that the mean of the stock returns is higher than the risk-free return r^b . Suppose that the individual has one unit of consumption to split between stocks and bonds. Let $u(x) = x^{1-\gamma}/(1-\gamma)$. Define

$$(A7) \quad g(s, \delta) \equiv E\{u'(1 + y\delta + sr^s + (1-s)r^b)(r^s - r^b)\}.$$

Then the partial derivative of g with respect to δ is

$$(A8) \quad g_2(s, \delta) = (-\gamma)E\{u'(\cdot)(r^s - r^b)(1 + y\delta + sr^s + (1-s)r^b)^{-1}y\}.$$

[Here the (\cdot) represents $(1 + y\delta + sr^s + (1-s)r^b)$.] Suppose $s^*(\delta)$ is defined so that $g(s^*(\delta), \delta) = 0$. Because y is independent of r^s and $E(r^s) > r^b$, it is clear that $s^*(\delta) > 0$. Our goal is to show that

$$(A9) \quad g_2(s^*(\delta), \delta) > 0.$$

If this is true, then standard comparative statics implies that $s^{**}(\delta) > 0$, which means that people with higher labor income invest more in stocks.

To prove that $g_2(s^*(\delta), \delta) > 0$ for a given δ , we impose the following sufficient condition:

$$(A10) \quad E\{\text{cov}(u'(1 + y\delta + s^*(\delta)r^s + (1-s^*(\delta))r^b), y(1+y+r^b)^{-1}|r^s)(r^s - r^b)\} < 0.$$

To gain some intuition into this condition, rewrite it as follows:

$$(A11) \quad 0 > E\{\text{cov}(u'(\cdot), y(1+y+r^b)^{-1}|r^s)(r^s - r^b)|r^s \geq r^b\} \times \text{prob}(r^s \geq r^b) \\ + E\{\text{cov}(u'(\cdot), y(1+y+r^b)^{-1}|r^s)(r^s - r^b)|r^s < r^b\} \times \text{prob}(r^s < r^b).$$

The concavity of u guarantees that the conditional covariance is always negative [because u' is decreasing in y while $y(1+y+r^b)^{-1}$ is increasing in y]. Hence, the first term on the right side of (A11) is positive while the second term is negative. If the conditional covariance were independent of r^s (as it would be if u' were linear), then (A11) would always be satisfied [because $E(r^s) > r^b$]. In fact, though, the third derivative of u is positive, so the conditional covariance is more negative for low values of r^s than it is for high values of r^s . (To see this, differentiate the covariance with respect to r^s , and note that u'' is increasing in y .) This raises the possibility that if u' is sufficiently convex and $E(r^s) - r^b$ is sufficiently small, then (A11) might fail. Note, though, that (A11) is only a sufficient, not a necessary, condition to prove that $g_2(s^*(\delta), \delta) > 0$.

Now consider the following chain of inequalities for arbitrary $s > 0$ and $\delta > 0$:

$$(A12) \quad E\{u'(\cdot)(r^s - r^b)(1 + y\delta + sr^s + (1-s)r^b)^{-1}y|y = \bar{y}\} \\ = E\{u'(\cdot)(r^s - r^b)(1 + y\delta + sr^s + (1-s)r^b)^{-1}y|y = \bar{y}, \\ r^s \geq r^b\}\text{prob}(r^s \geq r^b) \\ + E\{u'(\cdot)(r^s - r^b)(1 + y\delta + sr^s + (1-s)r^b)^{-1}y|y = \bar{y}, \\ r^s < r^b\}\text{prob}(r^s < r^b) \\ < E\{u'(\cdot)(r^s - r^b)(1 + y\delta + r^b)^{-1}y|y = \bar{y}, r^s \geq r^b\} \times \\ \text{prob}(r^s \geq r^b) \\ + E\{u'(\cdot)(r^s - r^b)(1 + y\delta + r^b)^{-1}y|y = \bar{y}, r^s < r^b\} \times \\ \text{prob}(r^s < r^b) \\ = (1 + \bar{y}\delta + r^b)^{-1}\bar{y}E\{u'(\cdot)(r^s - r^b)|y = \bar{y}\}.$$

Integrating over \bar{y} , we conclude that

$$(A13) \quad g_2(s, \delta)/(-\gamma) = E\{u'(\cdot)(r^s - r^b)(1 + y\delta + sr^s + (1-s)r^b)^{-1}y\} \\ < E\{y(1+y+r^b)^{-1}u'(\cdot)(r^s - r^b)\}.$$

This last term can be rewritten (using the independence of y and r^s) as

$$(A14) \quad E\{y(1+y+r^b)^{-1}u'(\cdot)(r^s - r^b)\} \\ = E\{E\{y(1+y+r^b)^{-1}u'(\cdot)(r^s - r^b)|r^s\}\} \\ = E\{E\{y(1+y+r^b)^{-1}|r^s\}E\{u'(\cdot)(r^s - r^b)|r^s\} \\ + \text{cov}(y(1+y+r^b)^{-1}, u'(\cdot)|r^s)(r^s - r^b)\} \\ = E\{y(1+y+r^b)^{-1}\}E\{u'(\cdot)(r^s - r^b)\} \\ + E\{\text{cov}(y(1+y+r^b)^{-1}, u'(\cdot)|r^s)(r^s - r^b)\}.$$

If $s = s^*(\delta)$, then the first term is zero (from the definition of s^*) and (A11) implies that the second term is negative. Hence, we can conclude that

$$(A15) \quad g_2(s^*(\delta), \delta)/(-\gamma) < E\{y(1+y+r^b)^{-1}\}E\{u'(\cdot)(r^s - r^b)\} = 0$$

which proves our theorem. Q.E.D.

Numerically

Finally, we demonstrate numerically the intuition about bond holdings substituting for risky labor income.

We follow Heaton and Lucas (1996) and assume that the household's income process is a two-state Markov chain with realizations 1.25 and 0.75, and the probability of exiting from one state to another is 0.26. As Heaton and Lucas do, we model stock returns as being independent and identically distributed over time with two equally likely realizations, 1.31 and 0.87. Stock returns are treated as independent of the income process. We assume that the real return to bonds is constant at 0.6 percent. We assume that the household lives for three periods and has a coefficient of relative risk aversion equal to 5. Its initial level of income is drawn from the stationary distribution of the income process.

For a wide variety of initial conditions of wealth (ranging from 0.01 to 100), we simulated 1,000 different sample paths of return and income realizations. For each of these samples, households reduced the portion of their wealth that they hold in stocks between the first and second periods. Intuitively, we know that the value of their human capital falls dramatically between these periods (because the number of remaining salary payments falls from two to one). Hence, households always compensate by buying more bonds.

References

- Bodie, Zvi; Merton, Robert C.; and Samuelson, William F. 1992. Labor supply flexibility and portfolio choice in a life cycle model. *Journal of Economic Dynamics and Control* 16 (July/October): 427–49.
- Fama, Eugene F., and French, Kenneth R. 1988. Permanent and temporary components of stock prices. *Journal of Political Economy* 96 (February): 246–73.
- Heaton, John, and Lucas, Deborah J. 1996. Market frictions, savings behavior, and portfolio choice. Department of Finance Working Paper 212. Kellogg Graduate School of Management, Northwestern University.
- Hodrick, Robert J. 1992. Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies* 5: 357–86.
- Ibbotson Associates. 1992. *Stocks, bonds, bills, and inflation—1992 yearbook*. Chicago: Ibbotson Associates.
- Jagannathan, Ravi, and Wang, Zhenyu. 1996. The conditional CAPM and the cross-section of expected returns. *Journal of Finance* 51 (March): 3–53.
- Kocherlakota, Narayana R. 1996. The equity premium: It's still a puzzle. *Journal of Economic Literature* 34 (March): 42–71.
- Malkiel, Burton G. 1996. *A Random Walk Down Wall Street: Including a Life-Cycle Guide to Personal Investing*. 6th ed. New York: Norton.
- Merton, Robert C. 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3 (December): 373–413.
- Morris, Kenneth M.; Siegel, Alan M.; and Morris, Virginia B. 1995. *The Wall Street Journal Guide to Planning Your Financial Future*. New York: Lightbulb Press.
- Quinn, Jane Bryant. 1991. *Making the Most of Your Money: Smart Ways to Create Wealth and Plan Your Finances in the '90s*. New York: Simon and Schuster.
- Ríos-Rull, José-Víctor. 1994. On the quantitative importance of market completeness. *Journal of Monetary Economics* 34 (December): 463–96.

Understanding Risk Aversion

How do different values of the coefficient of relative risk aversion lead to different implications for investor behavior? To get a feel for this, consider the accompanying table. It shows the risk premium that investors with coefficients of relative risk aversion γ demand in order to be indifferent between investing all of their wealth in a stock with a log-normal return and all of their wealth in a risk-free bond. (This assumes that the stock's log return has a standard deviation of 21 percent. Note that the risk premium demanded is approximately $0.5\gamma(0.21)^2$.) In the accompanying article, we often focus on the behavior of investors who have a risk aversion coefficient of 5.

The Extra Return Investors Need to Be Indifferent
Between Stocks and Bonds

Risk Aversion Coefficient (γ)	Risk Premium (% points)
.5	1.1
1.0	2.2
2.0	4.5
3.0	6.8
4.0	9.2
5.0	11.7
6.0	14.1
7.0	16.7
8.0	19.3
9.0	22.0
10.0	24.7

Table 1

The Horizon Doesn't Matter Much to Many Investors

Percentage of Wealth Invested in Stocks by Households With Various Degrees of Constant Relative Risk Aversion and Various Investment Horizons and No Possibility of Rebalancing Portfolios

Horizon (T)	Risk Aversion Coefficient (γ)					
	2	3	4	5	6	7
1	98.14	65.54	49.04	39.14	32.55	27.86
2	98.18	65.67	49.03	39.05	32.42	27.71
3	98.21	65.79	49.03	38.97	32.30	27.57
4	98.25	65.91	49.02	38.89	32.19	27.45
5	98.28	66.02	49.01	38.82	32.10	27.34
10	98.44	66.50	48.98	38.54	31.70	26.89
20	98.72	67.21	48.94	38.16	31.18	26.33
30	98.95	67.74	48.97	37.97	30.90	26.01
40	99.13	68.08	48.93	37.93	30.93	26.11

Table 2

A Measure of the Fear of Downside Risk

Minimum Amount Needed to Be Invested in a Stock Portfolio
for That to Be Preferred to \$1 Invested in an Optimal
Nonrebalancing Portfolio or a Bond Portfolio

Horizon (T)	Risk Aversion Coefficient (γ)			
	3		5	
	Optimal	Bond	Optimal	Bond
1	\$1.01	\$.981	\$1.04	\$1.02
5	1.04	.909	1.21	1.12
10	1.07	.827	1.45	1.25
20	1.14	.684	2.07	1.55
30	1.21	.566	2.89	1.90
40	1.27	.468	3.86	2.22

Chart 1

The Longer the Horizon, the More Likely That Stocks Will Outperform Bonds

Based on a Statistical Model Incorporating Data on the S&P 500 and U.S. Treasury Bills During 1926–90 and the Random Walk and Rational Expectations Hypotheses

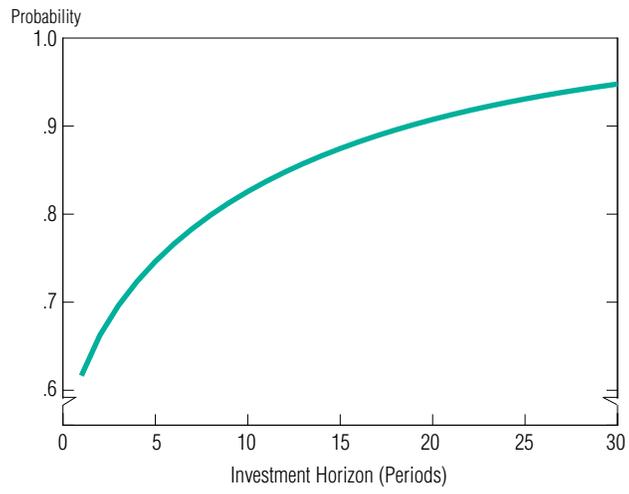
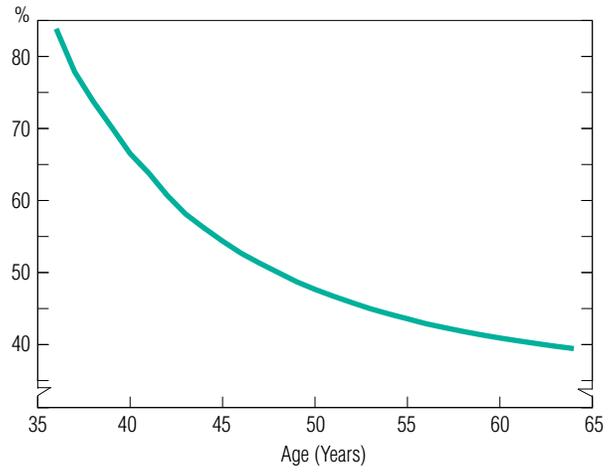


Chart 2

How a Typical Household's Stock Holdings Should Change Over Time

Percentage of Wealth Invested in Stocks by a Median Household, by Age, for a Household With an Annual Investment of \$20,000 From a Salary and a Risk Aversion Coefficient of 5*



*This is the median path of stock holdings from 1,000 randomly drawn time paths of stock returns.

How the Target Size Affects the Probability of Investing Less in Stocks

