Post-Retirement Financial Strategies: Forecasts and Valuation[1]

William F. Sharpe

Professor Emeritus, Stanford Graduate School of Business, Stanford, CA.

Introduction

This paper is concerned with financial strategies regarding the use of accumulated funds by an individual or couple after retirement. More specifically, I consider an investor or pair of investors with a fixed number of dollars to be used to finance consumption and possible bequests over a finite period of future years.

My goal is to develop a framework for evaluating a proposed strategy of this type, then use it to analyze a set of generic approaches based on industry practice.

To keep the analysis simple, I do not take into account other sources of income, liabilities, real property, insurance, taxes and many other important aspects that should be considered in many cases when adopting a post-retirement financial strategy. In practice, the issues considered here should be only part of the required analysis. But they would be an important part and I hope that this paper will contribute to the understanding of the associated issues by both academics and practitioners.

In most of the paper I follow the approach used in many papers in practitioner journals, assuming a fixed period of 30 years over which payments are made, without regard for mortality or the recipients of the fund (the initial investor or investors or beneficiaries thereof). For simplicity I will refer to the owner of the funds at retirement as "the investor", although the analysis would apply as well to cases in which the funds were those contributed by a couple.

In the latter part of the paper I argue that any of the strategies considered could be used as the basis for an annuity offered to a pool of investors with similar projected mortality probabilities. An annuity providing the same cash flows could be offered for a lower initial investment, reflecting the lack of any bequest after the annuitants' death or deaths. Alternatively, for the same initial outlay, larger payments could be provided. Importantly, most of the key aspects of the strategies apply whether an annuity is used or not.

---

[1] Keynote address at the European Financial Management Association Annual Meeting, June 24, 2011, Braga, Portugal.

The Capital Market Model

Financial advisors often advocate strategies in which funds are invested in multiple asset classes, such as domestic stocks, international stocks, bonds, commodities, etc. Here, in order to focus on the key aspects of financial strategies and to simplify the task of specifying the nature of capital markets, I utilize only two types of investments – one risky, the other riskless. Since investors are typically concerned with the real (inflation-adjusted) value of consumption, I measure all returns in real terms.

The riskless assets are inflation-protected treasury securities, which I will call TIPS, reflecting the terminology used for the Treasury Inflation Protected Securities in the United States. I assume that zero-coupon TIPS are available for every future year in the planning horizon. For simplicity I assume a constant term structure with a riskless real rate of 1% per year.

The risky asset is a “market portfolio” that includes all reasonably liquid bonds and stocks across the globe, each in proportion to its outstanding shares. This is inspired by asset pricing models (such as the Capital Asset Pricing Model and pricing kernel models) which imply that only the risk associated with the broad market will be priced and hence that efficient strategies should utilize a riskless asset and a broad market index fund. I assume that the expected one-year real return on the market portfolio is 4.5% and that its standard deviation is 10.0%. This provides an excess real return of 3.5% over that of the riskless security, giving a Sharpe Ratio (of the expected excess return to its standard deviation) of 0.35, broadly consistent with typical assumptions for highly diversified portfolios.

Throughout, returns are real and measured in value-relative terms. Thus a one-period return of 1.01 indicates that the real value at the end of the period is 1% greater than the value at the beginning. The one-period return on the market portfolio in period t will be denoted r_t. The cumulative return from period 1 through period t, V_t, will thus be:

$$V_t = r_1 \cdot r_2 \cdot \ldots \cdot r_t$$

Since the cumulative return equals the product of the periodic returns, it follows that the logarithm of the cumulative return will equal the sum of the logarithms of the periodic returns:

$$\ln(V_t) = \ln(r_1) + \ln(r_2) + \ldots + \ln(r_t)$$

I will make the relatively common, although not universal, assumption that market returns are identically and independently distributed (*iid*). In the first period a return is drawn from a specified distribution, in the next, a return is drawn independently from an identical distribution, and so on. In this sense, the process has no memory.

This is not an innocuous assumption. Some strategies for spending over time may well be motivated by a belief that returns are negatively serially correlated (equivalently, that security markets tend to revert to a long-run mean trend).

It remains to specify the particular distribution of returns for each period. Many practitioners assume that annual portfolio returns are normally distributed. But if a strategy's returns are independently and normally distributed for each period, they cannot be normally distributed for a horizon of two or more periods. Instead, the longer the holding period, the closer the long-term returns will conform to a log-normal distribution. This can be seen directly from equation (2), recalling that the central limit theorem states that the sum of a sufficiently large number of independent random variables, each with finite mean and variance, will be approximately normally distributed.

For simplicity, I will assume that the annual return on the market portfolio is lognormally distributed. Given independence, this implies directly that the return for any multi-year holding period will also be log-normally distributed. Equivalently, the logarithm of any single or multi-year return will be drawn from a normal distribution[2].

In a number of respects, this capital market model can be considered a discrete-time version of the continuous-time approach developed in [Merton 1971].

Pricing

Some post-retirement financial strategies can be evaluated using only assumptions about a multi-period return process such as the one I have described above. However, to estimate the costs of alternative strategies, it is necessary to have estimates of prices of possible future cash flows.

Consistent with the focus on strategies that utilize only the market portfolio and a riskless asset, I assume that only market returns are priced, both in any single period and also for any multi-period horizon. This could be consistent with a model of multi-period equilibrium in the capital markets, although I have no aspirations to develop one here (or elsewhere). In any event, as I will show, the assumption greatly restricts the characteristics of the pricing function.

---

[2] The following code (written in Matlab, the language used for all the computations in this paper) shows the relationship between the annual expected total return and standard deviation of return and the corresponding parameters for the moments of the logarithms of returns:

```
S = 0.10;
E = 1.045;
s = sqrt(log((S^2/(E^2)) + 1));
e = 0.5 *log((E^2)/exp(s^2));
```

Building on the approach of [Arrow 1952] and [Debreu 1959], I consider first a one-period setting in which there are *n* possible states of the world. Assume that securities are priced based on their payoffs in the states of the world, using a set of state prices, where p_s is the price today of $1 if and only if state s occurs at the end of the period. If the probability of state s is π_s, define the *pricing kernel value* for state s at time t as $m_{st} \equiv p_{st} / \pi_{st}$, a value that I called in [Sharpe, 2007a] the *price per chance* (PPC). The set of n values of m_{st} is the pricing kernel for time t. To say that only the market portfolio is priced in a single period is to assert that all states with the same market return have the same price per chance. Moreover, societal risk-aversion implies that the higher a state's market return, the lower should be its price per chance. More fundamentally, if markets are to clear, prices must adjust so that income in a state of scarcity (low market return) costs more than income in a state of plenty (high market return).

Now, consider a two-period case. Letting $\mathbf{m}_t$ and $\mathbf{r}_t$ represent vectors of pricing kernel values and market total returns (value-relatives) respectively for time t, the pricing kernels for periods 1 and 2 can be written as:

$$\mathbf{m}_1 = f_1(\mathbf{r}_1)$$
$$\mathbf{m}_2 = f_2(\mathbf{r}_2)$$

The pricing kernel for a horizon that includes both periods 1 and 2 will be the (dot) product of the two kernels. Thus:

$$\mathbf{m}_1 \cdot \mathbf{m}_2 = f_1(\mathbf{r}_1) \cdot f_2(\mathbf{r}_2)$$

In order for (1) the market to be priced in the same manner for each period and (2) for only the market to be priced for any multi-year horizon, it must be the case that:

$$f(\mathbf{r}_1) \cdot f(\mathbf{r}_2) = g(\mathbf{r}_1 \mathbf{r}_2)$$

A necessary and sufficient condition for this to be the case is that the one-period pricing function be isoelastic:

$$\mathbf{m}_t = \mathrm{A}\mathbf{r}_t^{-b}$$

More generally, if $\mathbf{M}_t$ represents the pricing kernel for payments t periods hence and $\mathbf{V}_t$ the cumulative market return over that horizon:

$$\mathbf{M}_t = \mathrm{A}^t \mathbf{V}_t^{-b}$$

Taking the logarithms of both sides of the equation:

$$\log(\mathbf{M}_t) = \log(A^t) - b \log(\mathbf{V}_t)$$

Clearly, the b coefficient indicates the elasticity of the pricing kernel with respect to cumulative market return – for every one percent increase in the latter, the pricing kernel decreases by approximately *b* percent. As is well known, this can be interpreted as indicating that a "representative investor" who holds the market portfolio has a utility function with a constant relative risk-aversion coefficient of *b*.[3]

The values of the coefficients in the pricing equation can be computed directly from the parameters of the assumed distributions of annual returns as follows[4]:

$$A = \left(\sqrt{E_m \cdot R_f}\right)^{b-1}$$

$$b = \frac{\ln(E_m / R_f)}{\ln(1 + S_m^2 / E_m^2)}$$

In this case:

$$A \approx 1.077$$
$$b \approx 3.74$$

I will use the pricing kernel primarily to estimate the values of the distributions of spending (also called "payments" or "paychecks") provided by a strategy). This requires a set of state prices – each of which represents the cost today of obtaining $1 at a given future time and state. Since the pricing kernel is simply a set of ratios of state prices to state probabilities:

$$\mathbf{P}_t = \mathbf{M}_t \cdot \Pi_t$$

Where $\mathbf{P}_t$ is a vector of state prices for payments at a future time t, and $\mathbf{M}_t$ and $\mathbf{\Pi}_t$ are, respectively, vectors for the pricing kernel and probabilities of the states for that time. In the simulations to follow, market returns were drawn randomly from the underlying probability distributions for *n* multi-year scenarios. Considering each scenario as a state, the probabilities all equal 1/*n* so that:

$$\mathbf{P}_t = \mathbf{M}_t / n$$

[3] For a further discussion, see [Sharpe 2007a].

[4] from [Scott, Sharpe and Watson 2009]. In the simulations to follow, to reduce sampling errors of draws from the underlying return distribution, the values of A and b used for each horizon were found iteratively based on the requirement that the implied present values of both the cumulative market returns and the cumulative risk-free return were both within a very small distance from 1.0. Since a million scenarios were generated for each simulation, the differences between the actual values of A and b and those that would have been obtained using the formula were small.

Investor Utility Functions

Many (but not all) economic models of multi-period consumption assume that the decision-maker can evaluate a strategy that provides consumption over one or more future periods by considering the probability distribution of consumption in each period separately from that in other periods, then summing the results. More precisely, the desirability of a set of uncertain amounts of consumption (here assumed to equal the payments made by a strategy) can be determined using a formula such as this:

$$EU = E[f_1(\tilde{x}_1)] + E[f_2(\tilde{x}_2)] + ... + E[f_n(\tilde{x}_n)]$$

Where:

$x_1, x_1, ..., x_n$ are the uncertain levels of consumption in periods 1,2,..n
$f_1, f_2, ... f_n$ are functions relating utility (desirability) to the amount received in periods 1,2,..n
E[ ] represents the expectation operator so that the utility of each outcome is weighted by its probability and the results summed
EU is the overall desirability or expected utility of the ranges of outcomes

In some axiomatic approaches built on the foundations provided by [Von Neumann and Morgenstern, 1944], the utility functions can be interpreted as equally desirable probabilities in a standard two-outcome gamble; in others, they are regarded as direct measures of a decision-makers' satisfaction with given outcomes.

The key attribute of such a measure is the fact that the expected utility of the outcome in each period can be determined without regard for outcomes that might have preceded it. Thus an investor with such preferences is said to have *time separable utility*.

In the current context, this implies that only the probability distributions of the possible outcomes in each of the future periods need be shown to the investor for him or her to evaluate the desirability of the strategy. And, if a choice is to be made between two strategies, comparison of their probability distributions for each of the future years should suffice. Moreover, for a strategy to maximize the investor's expected utility, the payments (x values) in a period must be a monotonic non-increasing function of the state prices and hence, in the capital market model used here, a monotonic non-decreasing function of the cumulative return on the market.

The time-separable utility model greatly simplifies the task of evaluating or designing post-retirement financial strategies. Given specific functions for an investor's utilities of consumption ($f_1, f_2, ... f_n$) and a set of state prices, one can easily create an optimal financial strategy for a complete market in which all needed state-contingent claims are available. As shown in [Sharpe 2006], this can provide a starting point for finding a desirable strategy in an actual market which does not include all such claims.

Unfortunately, both financial practice and behavioral research suggest that the time-separable utility model does not represent the preferences of all individuals or institutions. For at least some investors, the sequence of payments matters. At the most general level, such people wish to maximize the expected utility of the vector of payments:

$$EU = E[f\ (\tilde{x}_1, \tilde{x}_2, \ldots \tilde{x}_n)]$$

Even in a complete market, finding an optimal strategy for such an investor requires a specification of a utility function *f* and possibly a number of very complex calculations. The problem becomes even more difficult in an incomplete market.

To accommodate those who care about the sequence of payments over time in this manner, for each strategy analyzed I show two graphs summarizing probability distributions of future payments. One will be relevant for both those with time-separable utility functions and those with more general preferences; the other will provide additional information for those concerned with the sequence of payments. A third graph will provide estimates of costs and efficiency.

The Fidelity Income Replacement Funds

I now turn to the first of several possible post-retirement financial strategies, modeled on strategies offered by investment firms and financial advisors.

In 2007 Fidelity, one of the largest mutual funds in the United States, introduced a family of funds designed especially for use during retirement [Fidelity 2011]. Each such *Income Replacement Fund* has a target date at which the prototypical investor is assumed to have withdrawn all of his or her portion of the overall fund. In 2011, such funds were available with target dates ranging from 2016 through 2042.

A key characteristic of these funds is the availability of a payment program that can provide monthly payments each year based on the value of the assets in an investor's account at the beginning of the year. The "annual target payment rates" for an investment in a fund with thirty years remaining before its target date are shown in Figure 1[5].

[5] Source: [Fidelity 2011b]

Figure 1
Fidelity Income Replacement Fund Payment Rates

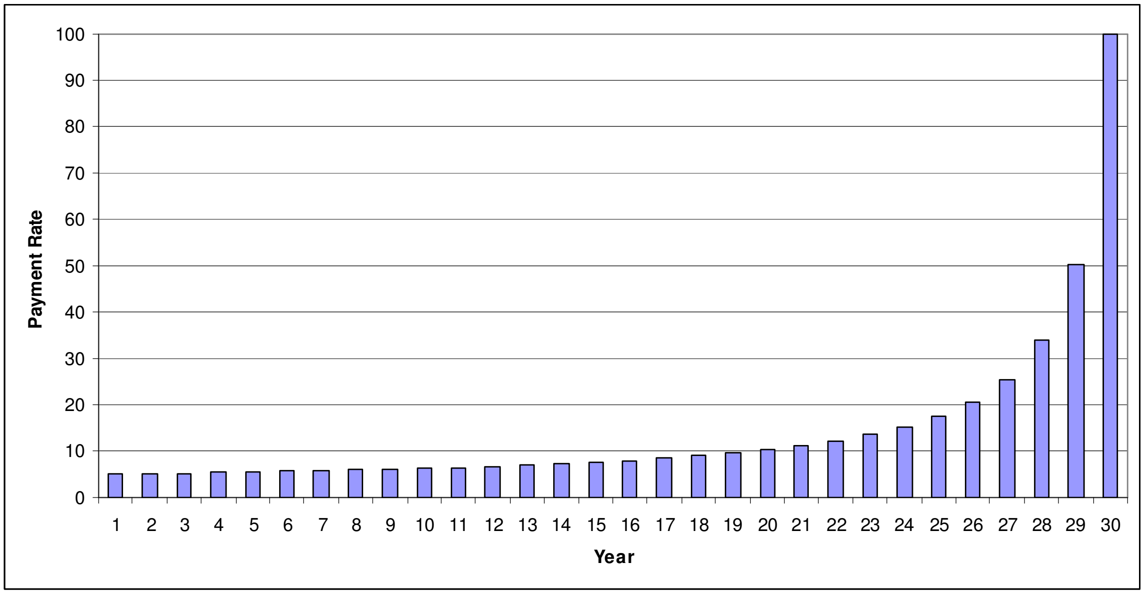

Each fund is invested in a number of underlying Fidelity funds, with the allocations varied through time to follow a "glide path" policy. Figure 2 shows the distribution among equity, bond and short-term funds for the 14 income replacement funds[6]. Detailed allocations among specific underlying funds are also shown in the fund prospectus.

[6] Source: [Fidelity 2011b]

Figure 2
Fidelity Income Replacement Fund Glide Paths

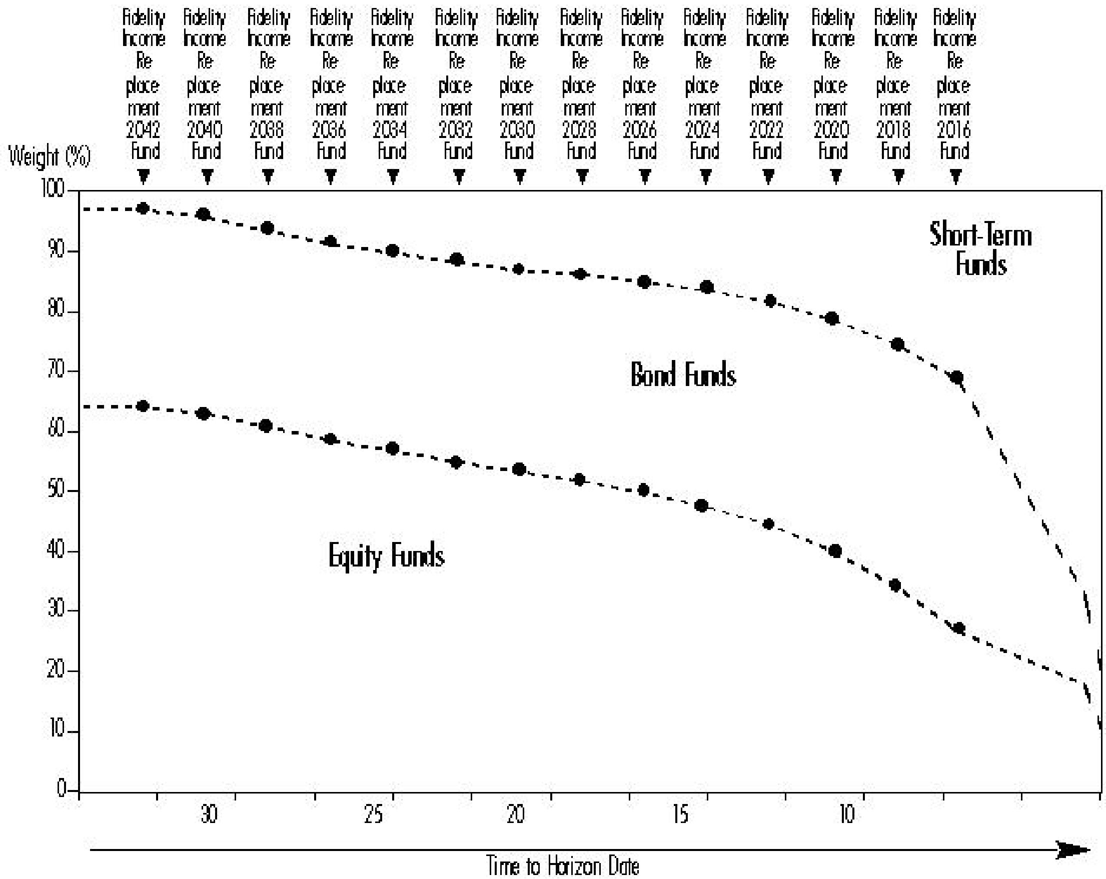

Fido

My goal here is to analyze a strategy (which I will call "Fido") that has the same payout rules as the Fidelity funds and a comparable glide path but invests only in a market portfolio and a riskless real asset. To do so, I first estimated the standard deviations of the bond and stock portions of Fidelity's portfolios using the representative parameters in the text provided by the CFA Institute[7], then divided the result by 10.0%, the standard deviation for the market portfolio that I have assumed. The resulting proportions invested in the market portfolio each year are shown in Figure 3.

[7] [Maginn, et. al., 20070]. The estimates, provided by UBS Global Asset Management, were: U.S. Fixed Income standard deviation = 5.7%, U.S. Equity standard deviation = 15.4%, correlation = 0.25,

Figure 3
Fido: Assumed Proportions Invested in the Market Portfolio

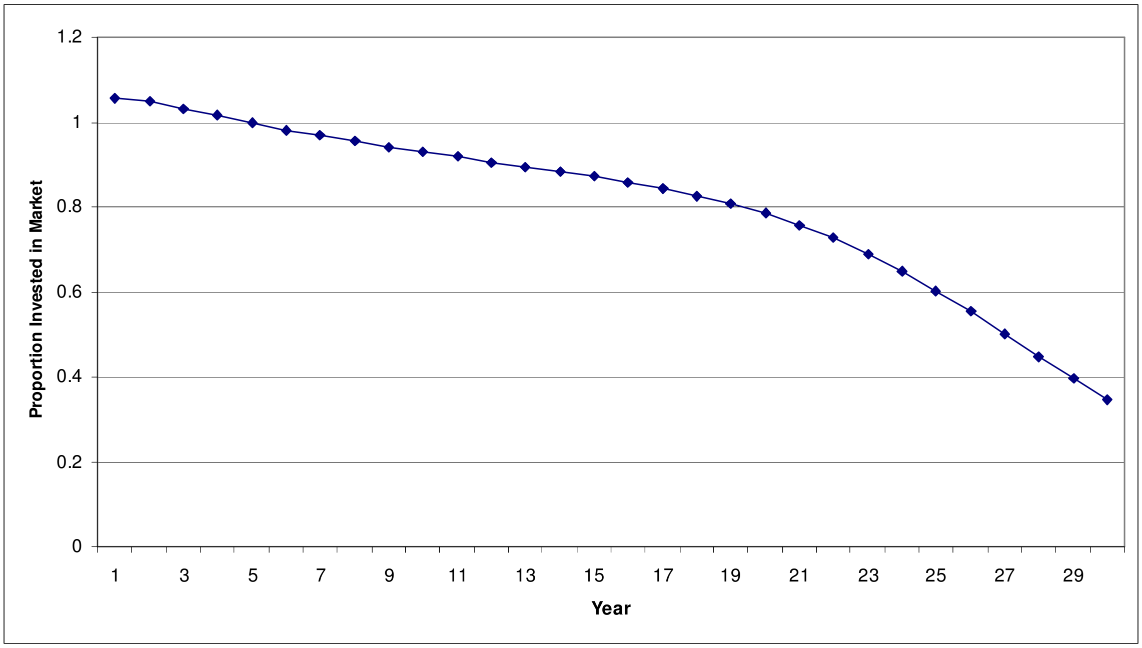

With these inputs and the assumed capital market return generating process it is straightforward to use a Monte Carlo analysis to simulate the payments generated by this or any other strategy. To minimize sampling error, I generated one million scenarios for each of the simulations in this paper. To insure comparability, for each strategy analyzed, I used the same set of random draws from a lognormal distribution to determine returns on the market portfolio over the course of 30 future years, although tests showed that the key results were not sensitive to the set of random variables utilized. Each strategy was assumed to be followed by an investor with an initial wealth of $1 million.

The distribution of the million possible payments in year 30 for the Fido strategy is shown in Figure 4. It is similar to the usual portrayal of a cumulative probability distribution but inverted so that the vertical axis shows the probability of exceeding any specified goal on the horizontal axis – a portrayal that provides more direct answers to questions often posed by individual investors.

Figure 4
Fido: Distribution of Payments in Year 30

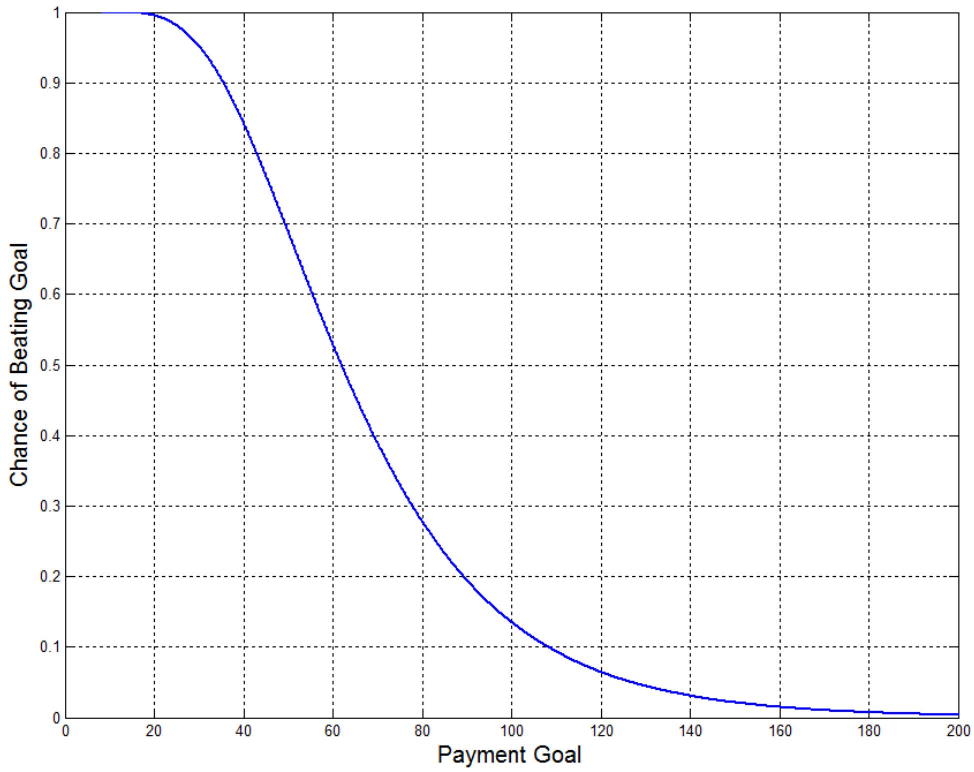

Unfortunately, a graph with detailed distributions of this form for each of 30 years provides too much information to be processed by a typical investor attempting to evaluate a particular retirement strategy, let alone compare alternative approaches. Instead, I summarize the salient information in the form shown in Figure 5. For each future year, I include bars shaded to show the goals which can be beaten with probabilities 0.99, 0.95, 0.75, 0.50, 0.25, 0.05 and 0.01. For added emphasis, the median outcomes are highlighted with a circle.

Figure 5
Fido: Payment Distributions

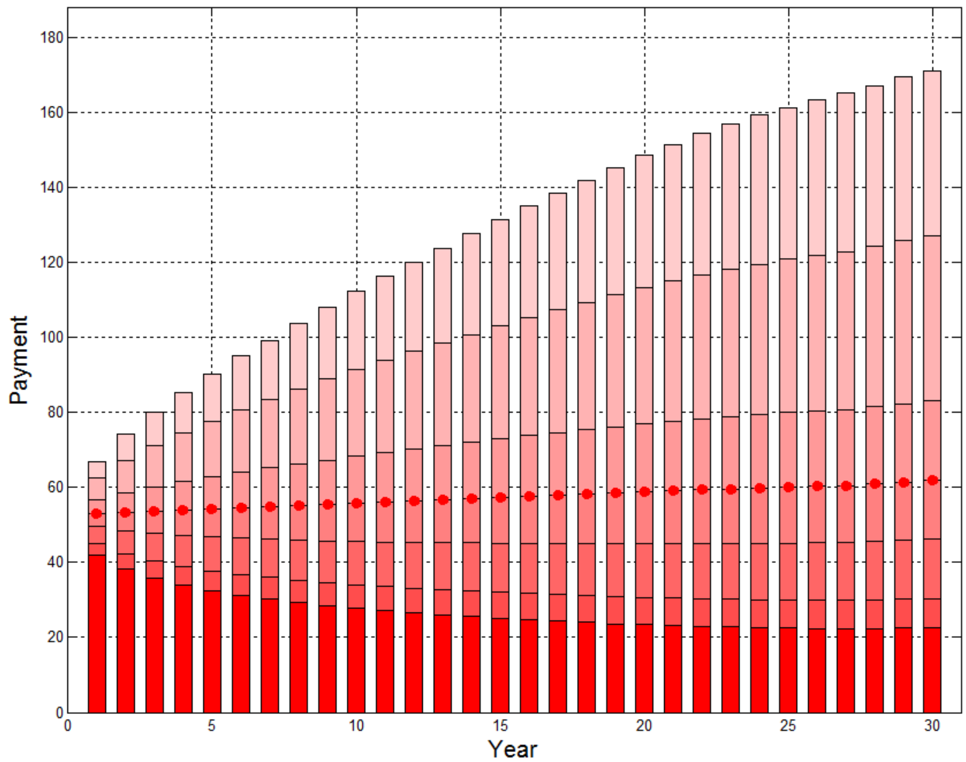

As can be seen, the median outcomes are slightly greater, the farther in the future the year of receipt. However, the range of possible payments becomes considerably greater, the farther in the future is the year the payments would be made. This may surprise some, since one appeal of a glide path strategy is the reduction in risk over time as more conservative asset mixes are utilized. But every year's payment is based on a compound return that includes results from earlier, more risky allocations. Accordingly, seen from today, there is considerably greater uncertainty about later payments than earlier ones.

Forecasts such as those shown in Figure 5 should prove useful for any investor who wishes to evaluate a post-retirement financial strategy. And for those with time-separable utility, the information summarized in this form may also be sufficient, since only the probability distributions of payments in each year are germane. While such investors may wish to inspect the 30 underlying distributions in full, no information about probable sequences of returns is needed.

For concerned as well with the sequence of payments a graph such as that shown in Figure 6 may also be helpful. Each bar summarizes the distribution of the ratios of the payment in a year to the payment in the prior year. In this case the median outcomes are all slightly greater than 1.0, reflecting an expectation of slight increases in payments over time. But there is considerable uncertainty. In this respect, however, the effect of the glide path is more evident. Viewed from the prior year, uncertainty about the payment in the next year is smaller in the later years than in the earlier ones.

Figure 6
Fido: Payment Ratio Distributions

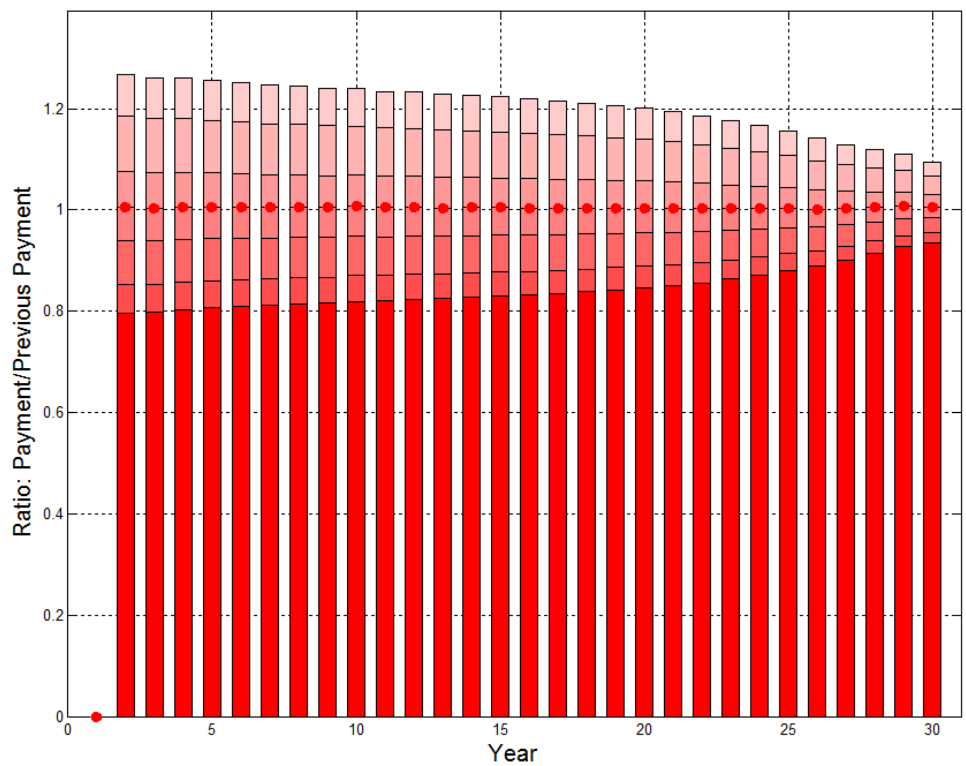

The addition of the Payment Ratio Distributions graph should prove helpful for those concerned with a strategy's sequence of payments. However, even if the full underlying distributions of these ratios were considered, an investor with a general utility function might still be able to only approximate the expected utility of a strategy. This said, the combination of a graph of the Payment Distributions (such as Figure 5) with that of the Payment Ratio Distributions (such as Figure 6) should prove useful information in a parsimonious and relatively easily understood manner.

Lockbox Strategies

In [Sharpe 2007b] I showed that in a complete market setting where derivatives or dynamic strategies can provide any desired set of payments over time and in different states of the world, any post-retirement financial strategy can be implemented by setting up a series of *lockboxes*, one for each future year. Each such box contains an initial amount of money and information on a specified investment strategy to be followed until its designated maturity year; at that point, the investments will be liquidated and the proceeds spent.

This result is of theoretical interest. However, in an incomplete market many proposed post-retirement financial strategies cannot be economically or even feasibly implemented in this manner. But some can, including Fidelity's Income Replacement Funds and the Fido strategy that uses their payment rules.

To see this, let $p_1, p_2, ..., p_{30}$ be the annual payout rates shown in Figure 1, expressed as proportions of the fund's current asset values. Let $\tilde{r}_1, \tilde{r}_2, ..., \tilde{r}_{30}$ be the annual returns on the investment strategy, expressed as value relatives, and w_O the initial amount to be invested. Then the first payment will equal:

$$\tilde{x}_1 = (w_0 \tilde{r}_1) p_1 = (w_0 p_1) \tilde{r}_1$$

which is equal to the terminal value of a lockbox with an initial value of $w_0 p_1$ invested for one year.

Similarly, the second payment will equal:

$$\tilde{x}_2 = \{w_0 \tilde{r}_1 (1 - p_1) \tilde{r}_2 p_2 = \{w_0 (1 - p_1) p_2\} \tilde{r}_1 \tilde{r}_2$$

which is equal to the terminal value of a lockbox with an initial value of $w_0 (1 - p_1) p_2$ invested for two years.

Extending this approach gives the proportions of the initial investment to be placed in each of the lockboxes. In this view, the Fidelity and Fido strategies can be considered lockbox strategies in which each lockbox uses the same glide path investment rule up to its maturity date. Equivalently, each of the initial values for the lockboxes represents the present value of the payment distribution for the associated maturity year.

Valuation

Thus far I have not utilized the pricing kernel developed earlier, since forecasts of payment distributions and payment ratio distributions can be made using only assumptions about the return generating process for the assets employed. Moreover, for some strategies, such as Fido, each year's payment distribution can be valued without making assumptions about state prices. But this is not the case for all strategies, and it is useful to estimate the values or costs of other alternatives. I will do so using the isoelastic kernel, with the caveat that the results will, at best, be estimates of actual values that would be observed in a competitive market.

Consider the payment distribution in year 30 for the Fido strategy. Based on the previous formulas it has a present value equal to 2.213% of the initial investment. Using the computed state prices, the present value is the same to three decimal places. However, the valuation function implies that this is not the least expensive way to obtain this probability distribution. Figures 7 and 8 (plotted using logarithmic scales due to the wide range of results) show why.

Figure 7
Fido, Year 30:
log(Payment) vs. log(Cumulative Market Return)

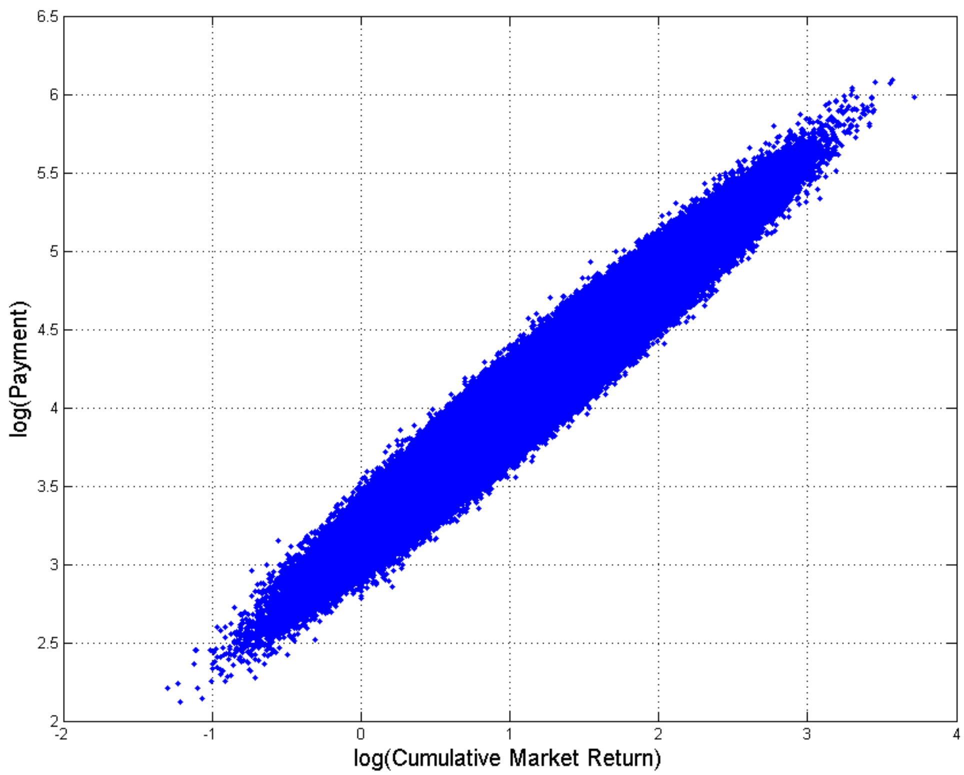

Figure 8
Fido, Year 30:
log(Price) vs. log(Payment)

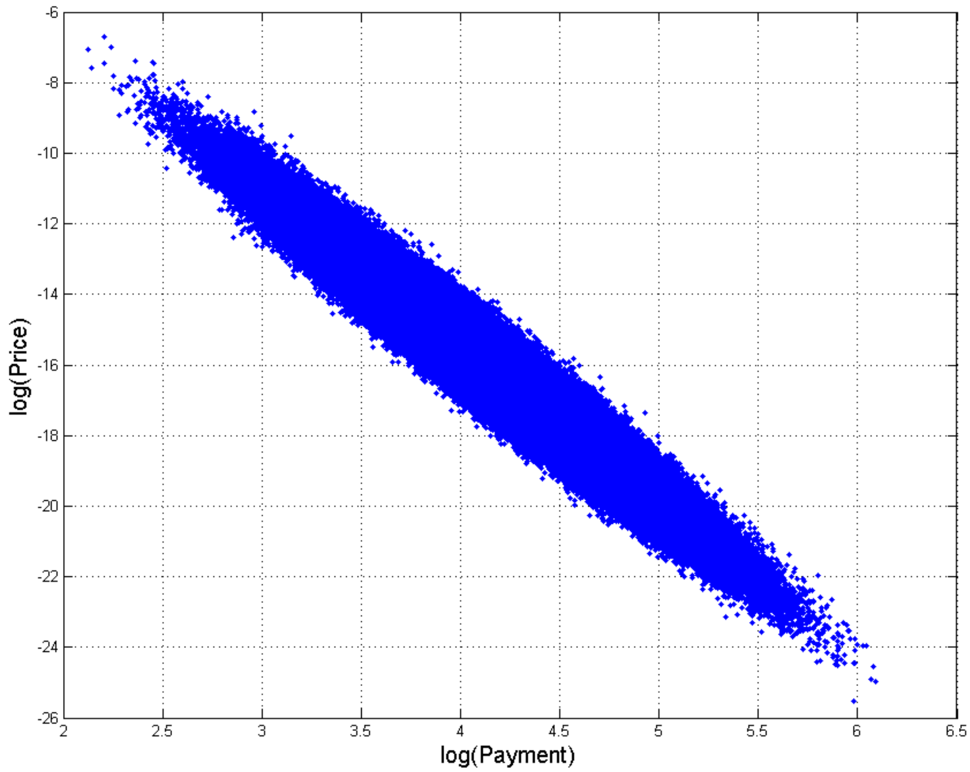

Figure 7 shows that the payments depend not only on the cumulative return on the market over the thirty years but also on the paths that returns take over the years. But the valuation model posits that state prices are a function only of cumulative market returns. Thus, as shown in Figure 8, payments are not a monotonic function of state prices[8]. As shown in [Dybvig 1988a] and [Dybvig 1988b], this, in turn, implies that there is a cheaper way to obtain the Fido probability distribution of payments in year 30.

To see why, let *p(x)* represent the state price for the state in which a payment *x* is obtained. Consider two payments, x_a and x_b, where $x_a > x_b$, the costs are respectively p_1 and p_2, and $p_1 > p_2$. The cost of the two payments will be $p_1\,x_a + p_2\,x_b$. Now, switch the two payments so that x_a is received the state with price p_2 and x_b in the state with price p_1. The new cost, $p_2\,x_a + p_1\,x_b$, will be lower since the greater payment is to be made in the cheaper state and the smaller in the more expensive state. Clearly, any allocation of payments to states in which there is a situation such as this can be altered to lower the overall cost.

---

[8] Figures 7 and 8 could be somewhat misleading since many more payments plot along the center of the relationship than on the periphery. In fact, the R-squared value for each of the relationships is 0.946.

It is a simple matter to find the lowest possible cost allocation for any set of payments by simply (1) sorting the vector of payments with the largest first, (2) sorting the vector of state prices with the smallest first then (3) assigning the payments in the first vector to the states in the second vector in order. The result will be a set of payments with the same probability distribution as the initial one and the lowest possible cost. Importantly, the sets of payments are the same; only their assignments to states differ. Unlike the original allocation of payments across states, they are *path-independent*.

Given our model in which state prices are a function only of cumulative market returns, path-independent payments will be a non-decreasing function of such returns, as shown in Figure 9. In the case at hand it is possible to provide a similar (although not identical) set of payments with a simple investment strategy (here, rebalancing the investments in a year 30 lockbox each year to a predetermined mix by asset value). However for many strategies this may not be the case.

Figure 9
Fido, Year 30:
log(Path-independent Payments) vs. log(Cumulative Market Return)

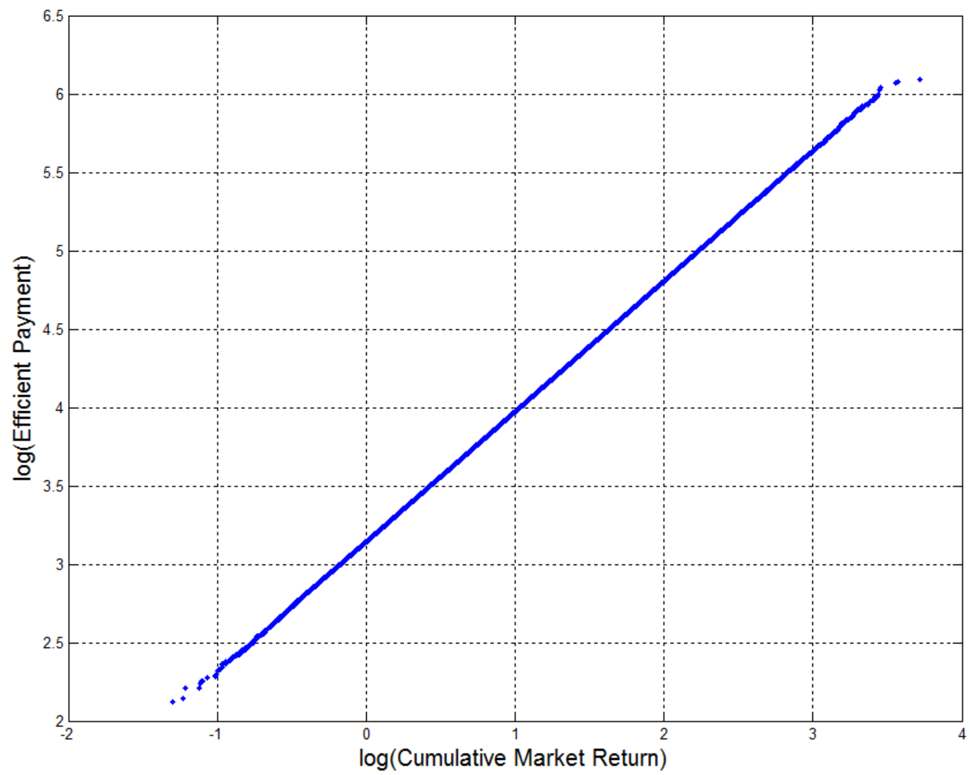

It remains to determine the present value of the path-independent payments. This requires only that they be priced by multiplying each payment by the state price for the state in which it is to be received, then summing the results. In this case the present value is 97.6% of the present value of the original set of payments.

For an investor with time-separable utility, the path-independent alternative is clearly preferable. The same probability distribution of payments in year 30 can be obtained for 97.6% of the cost of the original strategy. Alternatively, the cost of the outcomes from the original strategy could be used to obtain a path-independent set of payments that would provide more income in every possible state of the world. This suggests that for such an investor the glide path investment strategy would be dominated by the strategy shown in Figure 9.

For an investor with a more general utility function, the choice is less clear. If the payment distributions in all the years were changed to their path-independent counterparts, such an investor might consider the time-series behavior of returns less desirable. Path-independent counterparts for a path-dependent strategy may not be preferred by all.

Figure 10 shows the present values of the payments in each of the thirty years for the Fido strategy. The green bars represent the present values of the alternative path-independent outcomes, the sum of the green and magenta bars the present values of the original outcomes, and the light bars the additional cost due to path dependency. As can be seen, the additional cost is significant only for the later years. Overall, the sum of the present values for the path-independent alternatives is only slightly less than that for the original outcomes – the ratio of the former to the latter is 99.72%. Anyone who finds the year-to-year behavior of the payments from this very gradual glide path strategy preferable might well be willing to accept the small amount of inefficiency in obtaining the annual payments.

Figure 10
Fido: Present Values of Payments

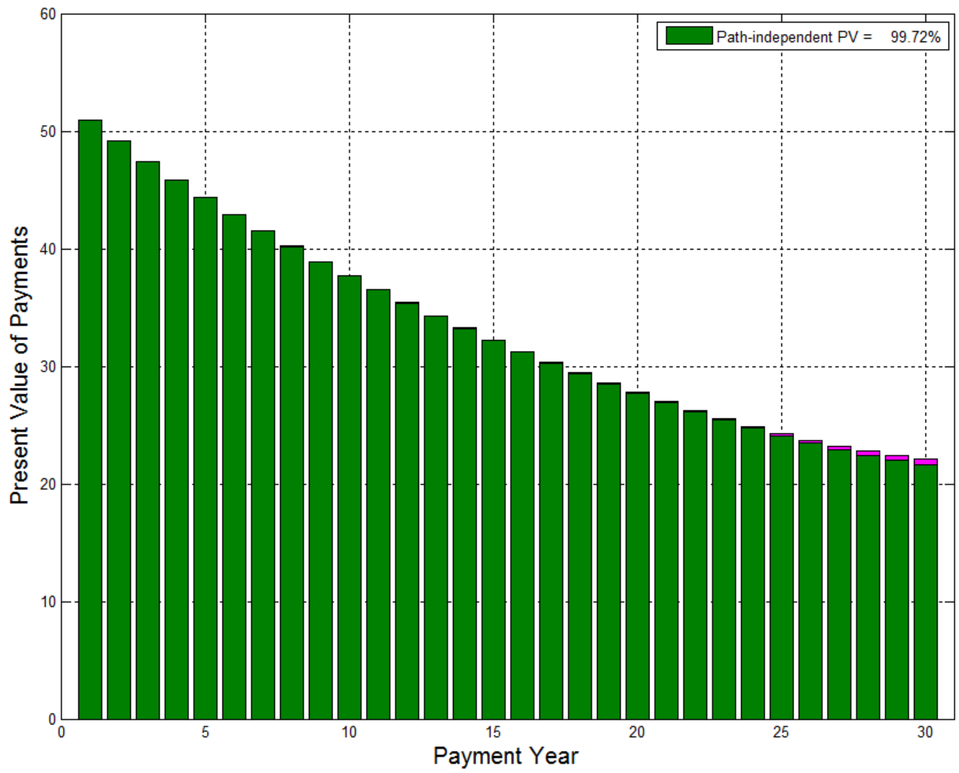

This completes the development of the tools that I propose be used to help evaluate post-retirement strategies. In the remainder of the paper, I utilize them to investigate the properties of three other approaches.

Vanguard Managed Payout Funds

Vanguard, another major mutual fund company in the United States, offers a set of three Managed Payout Funds. According to the prospectus [Vanguard 2010], the funds are "…primarily intended for retirees and other investors who desire monthly payments to meet expenses, who are able to tolerate the inherent risks of Managed Payout Funds, and who want to retain access to their investments to meet unexpected expenses or, potentially, for estate planning." The funds "…are designed, at your choice, either to be the sole fund in your portfolio or to complement other investments as part of your overall investment strategy, depending on your personal needs and preferences." Particularly germane for the purpose of this paper, the funds are said to be "…innovative in that they have a 'built in' systematic withdrawal plan (SWP)…" and that a "unique feature of Managed Payout Funds is their three-year rolling average payout strategy, which can lead to a more consistent payout level than typical SWPs."

In the United States, college and university endowments often use this sort of asset smoothing[9]. A recent study [NACUBO 2010] found that 75% of responding institutions spent a percentage of a moving average of asset values and only 4% spent a percentage of the market value of assets.[10] In this and other respects, the Vanguard funds are designed to follow the "endowment model".

Each of the three funds has a target mix of stocks and bonds and a pre-specified payout rate. The actual percentages in equity and payout rates in March, 2011 were those shown in Table 1.

Table 1
Vanguard Managed Payout Funds

Fund	% Stocks, 3/31/2011	Payout: % of 3-year Average Daily Assets
Distribution Focus	65.5 %	3 %
Growth and Distribution	75.3 %	5 %
Growth Focus	84.1 %	7 %

These funds are not designed for a specific holding period, as are the Fidelity Income Replacement Funds. I will not try to analyze a direct counterpart here. Instead, I will focus on the use of average of historic values in determining each year's payout.

Vido

The section analyzes a strategy that incorporates all the features of the Fido approach with one exception: each year's payout ratio is applied to the average up to four years' ending asset values – the current value and three prior years. For the first three years, all available year-end values are utilized; thereafter the averages of the current and three prior year's values are employed . As with the Fido strategy, in the final year the entire remaining value of the fund is paid out. Since the approach incorporates aspects of both the Fidelity and Vanguard funds, I call it Vido.

[9] Traditional defined-benefit pension funds often behave similarly; basing employer contributions on smoothed values of assets, liabilities and/or past benefit increases.
[10] The remaining 21% used some other basis

Figure 11 shows the resulting payment distributions. Notable is the final year, in which the payment equals the remaining assets. Such payments are considerably smaller, since in earlier years payments were based on prior and generally larger asset values than those of the actual fund in prior years.

Figure 11
Vido: Payment Distributions

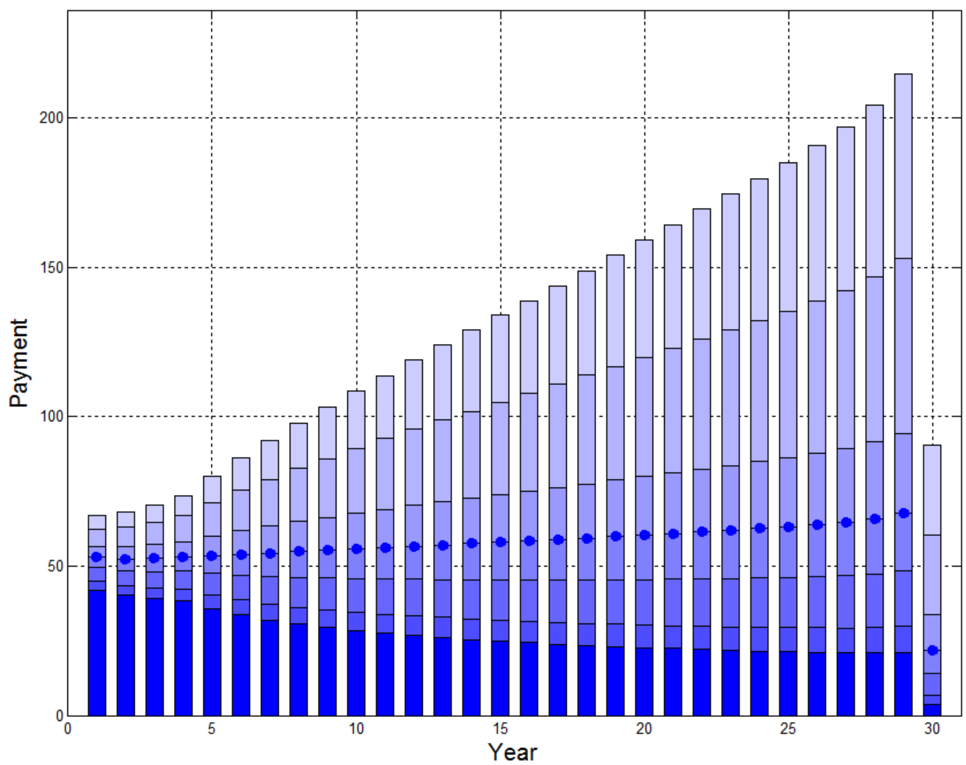

Figure 12 provides a comparison of the Fido and Vido payment distributions. For each year, the bar on the left shows the range of payments for the Fido strategy and the bar on the right the range of payments for the Vido strategy. In the early years after year 1 Vido's use of smoothed asset values reduces the range of payments, as intended. However, in the later years the effect is reversed. This illustrates a key point. Ultimately, the risk of a strategy depends on the manner in which the assets are invested. A payment (spending) rule can shift this risk from some years to others, but it cannot reduce or eliminate it.

Figure 12
Fido and Vido: Payment Distributions

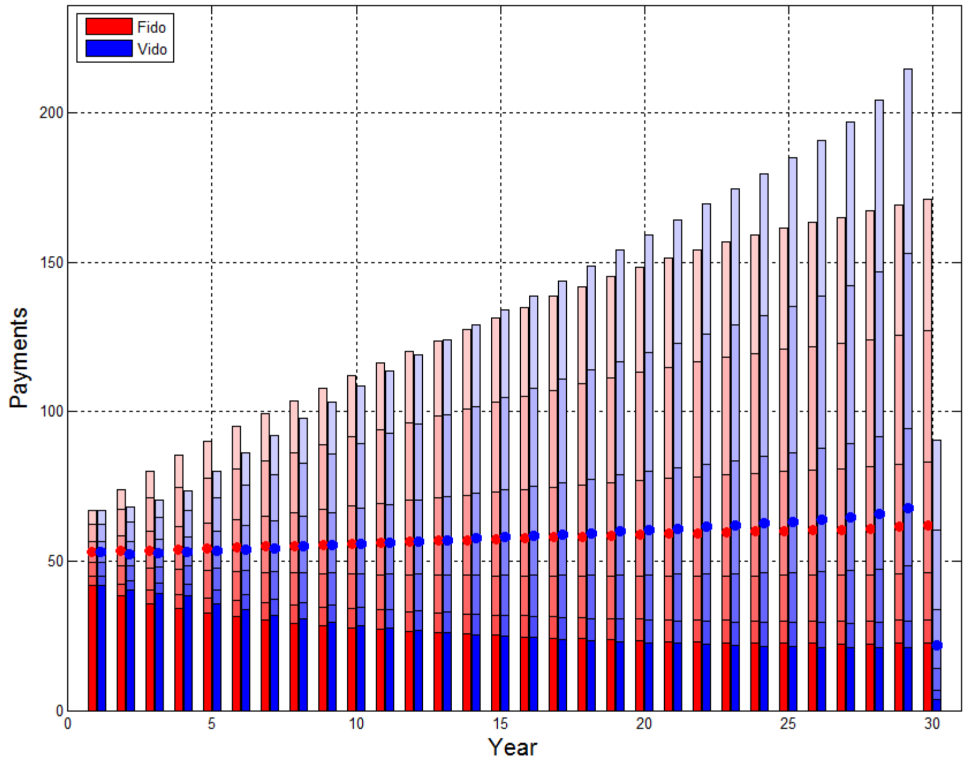

A similar relationship is shown in Figure 13, which compares the payment ratio distributions for the two strategies (again, with Fido on the left and Vido on the right for each year). Asset smoothing decreases the year-to-year variability of payments from years 2 through 27 but thereafter the variability is greater, and dramatically so in the final year.

Figure 13
Fido and Vido: Payment Ratio Distributions

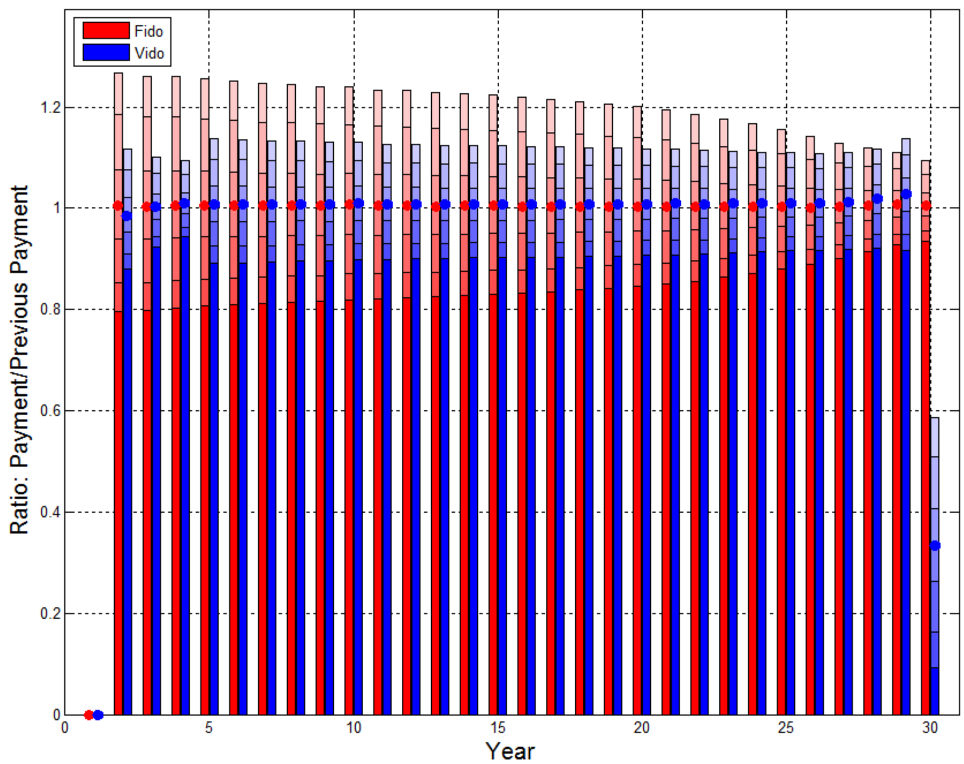

One may well ask why endowments and many other institutional funds so often choose to smooth their payouts, only partially taking into account changes in asset values. A plausible reason may be an explicit or implicit belief that asset values tend to be negatively serially correlated rather than independently and identically distributed, as assumed here. Other reasons may be the hoped-for perpetual life of the institutions, coupled with the ability to obtain new sources of funds when needed.

Figure 14 shows the present values of the payments for the Vido strategy. In this case all the computations are based on state prices since a simple lockbox strategy is not feasible. As with Fido, the present values of the path-independent strategies with the same annual payment distributions are smaller. As anticipated, basing payments on average historic asset values adds to the path dependency. However the overall dollar cost of the strategy's path-dependence is still relatively small – a path-independent strategy would cost 98.31% as much (compared with 99.72% for the Fido strategy).

Figure 14
Vido: Present Values of Payments

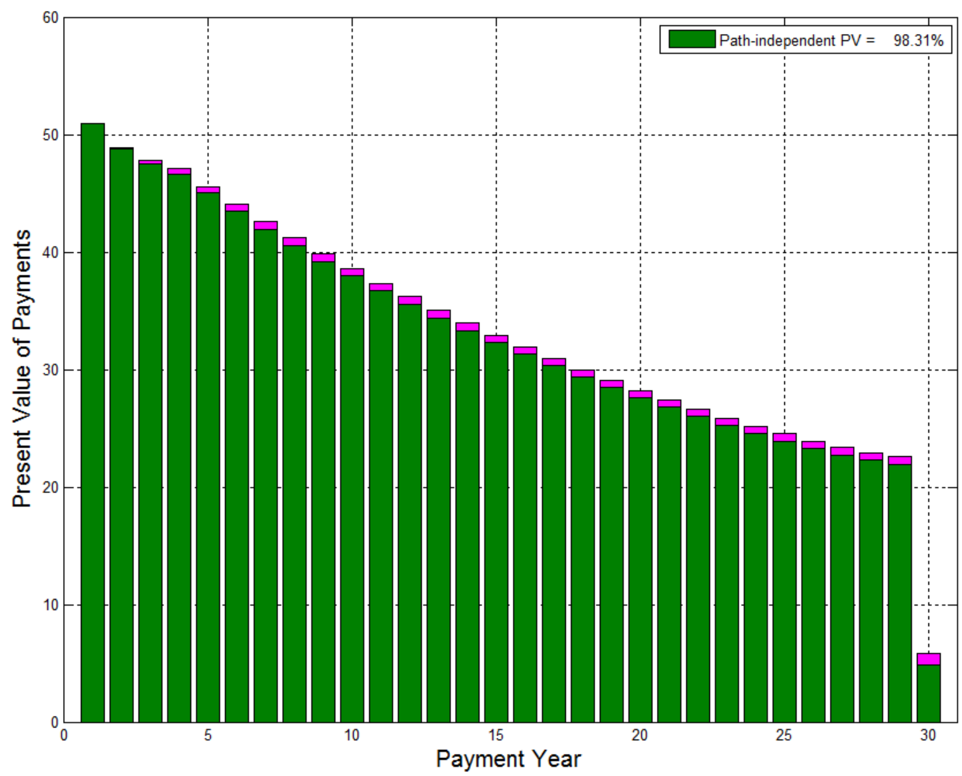

The Four Percent Rule[11]

Many Financial Advisors endorse a simple post-retirement investment and spending strategy: invest in a combination of stocks, bonds and cash with relatively high expected returns, spend an initial percentage of overall savings, then increase spending each year at the experienced rate of inflation. This approach, developed in [Bengen 1994], was derived by testing different amounts of initial spending using historic asset returns. After doing so, Bengen found that a policy of starting with spending equal to 4% of savings, then maintaining a constant amount of real spending would have provided a high probability that the investor did not "run out of money" in 30 years. The generic approach

[11] Much of the material in this and the next section parallels the analysis in [Scott, Sharpe and Watson 2009].

is now usually termed the “Four Percent Rule”, whether the initial spending is four percent of savings or some other amount.

A great many investment and advisory organizations suggest that retirees consider an approach of this sort. Figure 15 show a representative planning tool from the Vanguard web site[12]. The user provides an initial balance, chooses one of three asset allocations, and specifies the number of years to be “spent in retirement”. The tool then provides an initial withdrawal rate (per year) and the corresponding initial monthly withdrawal in dollars.

Figure 15
The Vanguard Retirement Planning Tool

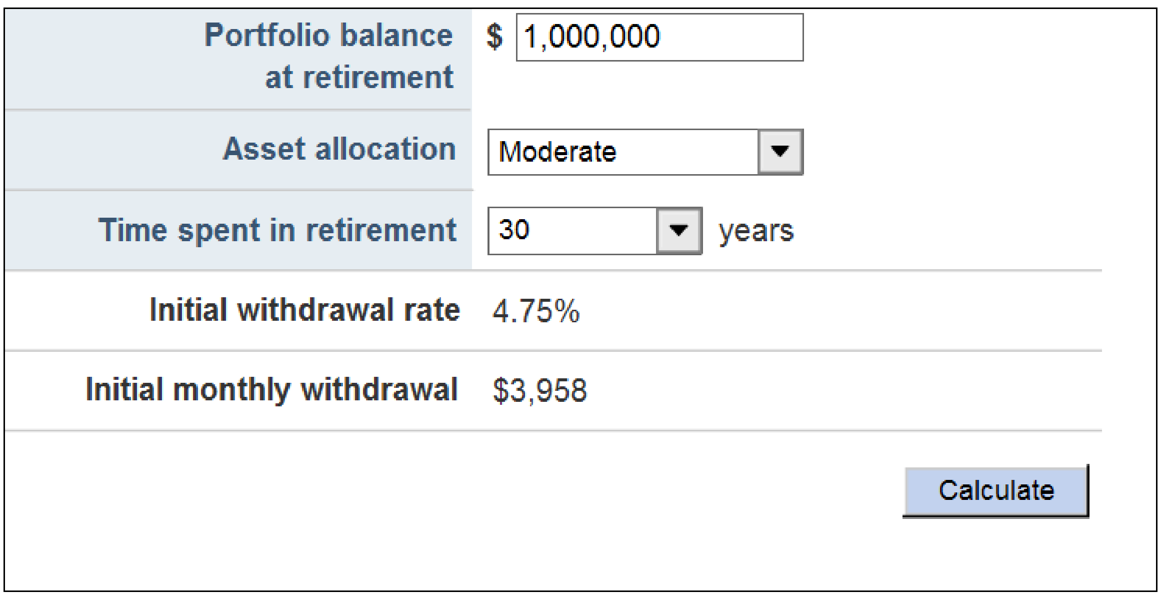

The investor is assumed to maintain a relatively constant asset allocation over the full period using one of the three portfolios:

> “…a conservative portfolio is one in which stocks make up no more than 35% of the mix, a moderate portfolio is 35% to 65% stocks, and an aggressive portfolio is at least 65% percent stocks.”[13]

I

[12] [Vanguard Withdrawal in Retirement tool, 2011]
[13] Ibid.

The basis for the tool's recommended initial spending is typical of the research in this area:

> "In this calculator, we create 83 different simulated "time paths" for the evolution of your portfolio by first assuming that you begin taking withdrawals at a specific point in history (e.g., 1926, 1927, 1928). We then use the actual, historical rate of return that occurred in each subsequent year from that point forward, applying them in sequence to your portfolio as time rolls forward. If such a path needs to go beyond the year 2008, we just loop back to the returns of 1926 and cycle forward from there … Given the past historical data we use (returns from 1926 through 2008), we end up with 83 different starting years with 83 different time paths. The monthly withdrawal amount shown by the tool is the highest level of spending in which 85% of these historical paths would have left you with a positive balance at the end of your chosen investment horizon."[14]

Cuatro

To explore some of the characteristics of this approach, I analyze Cuatro, a strategy designed to follow a typical implementation of the four percent rule. The first year's spending equals four percent of initial wealth. The goal is to make payments with the same real value for 30 years. Throughout the period, funds are invested solely in the market portfolio.

Figure 16 shows the payment distributions (in $thousands) for an investor with an initial wealth of $1,000,000. In each of the first 20 years there is at least a 99% chance of receiving the desired real payment of $40,000. In year 21 there is a 99% chance of receiving a real amount of only $29,000 or more. Thereafter, the probability of receiving the desired real value of $40,000 is less than 99% but more than 95%.

[14] Ibid.

Figure 16
Cuatro: Payment Distributions

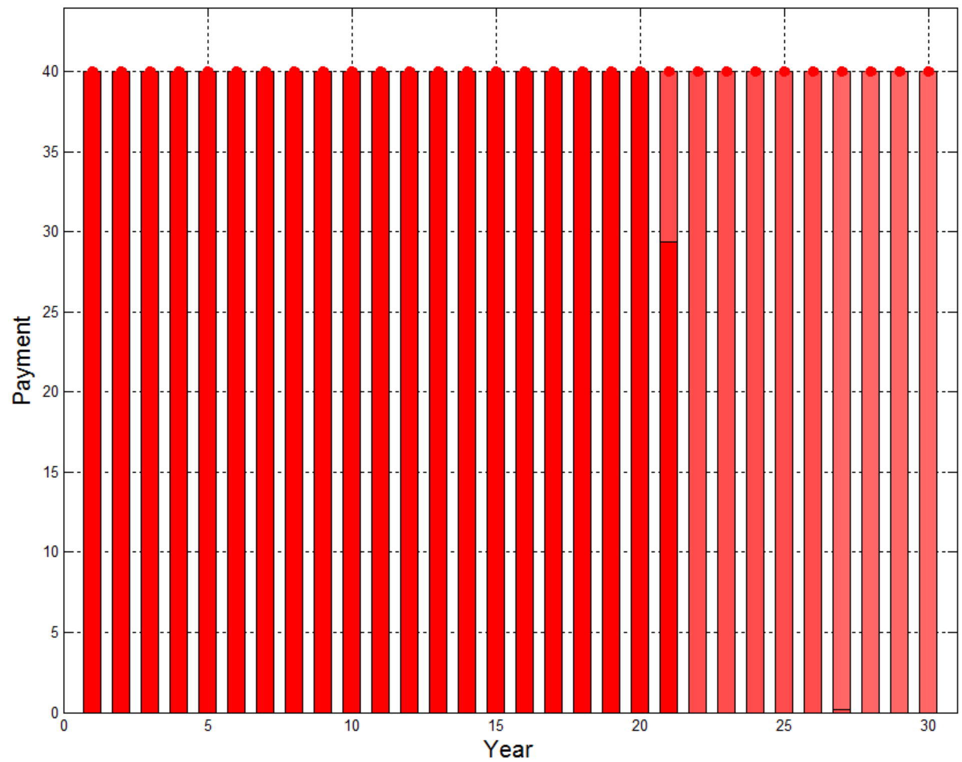

Figure 17 show the distributions of the ratios of year-to-year distributions. Clearly, the desire to obtain smooth payments from an investment with highly variable returns works only to a point. In the later years there is a chance that the payments will fall to zero. As with the Vido strategy, risky investments cannot provide riskless payments forever.

Figure 17
Cuatro: Payment Ratio Distributions

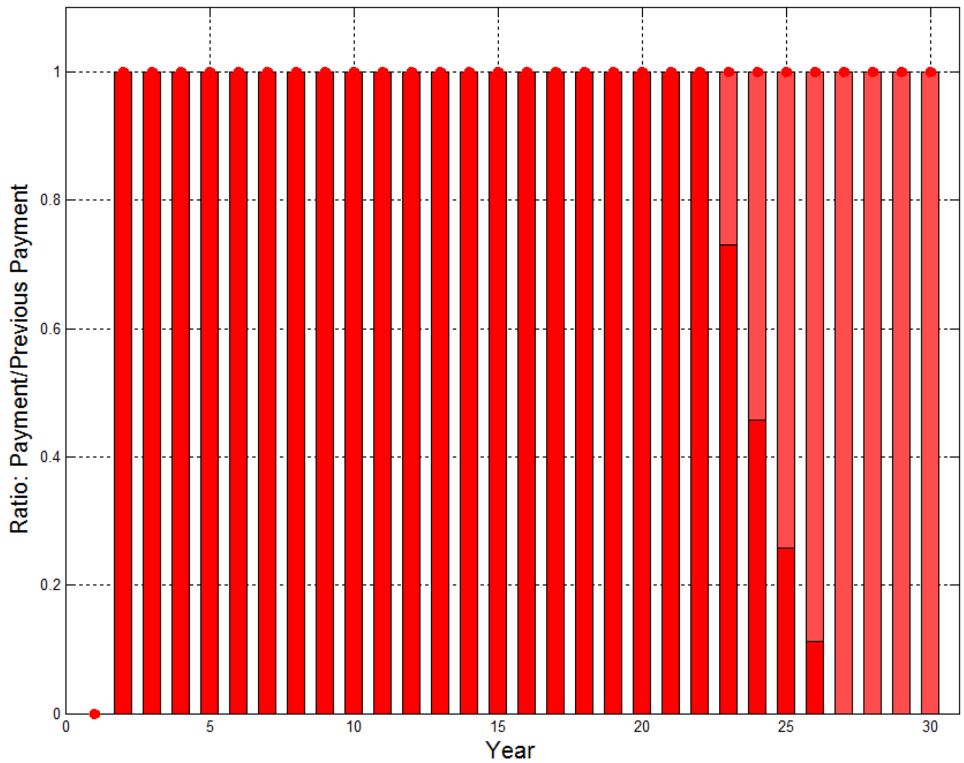

Figure 18 shows the present values of the payment distributions for the 30 years in the investor's horizon. They are smaller for later years due to discounting at the real interest rate. However, the values begin to decline at an increasing rate after year 16, reflecting the effects of scenarios in which payments are small or zero and the fact that such scenarios tend to be characterized by low cumulative market returns and correspondingly high state prices. Payments in the latter years become increasingly path-dependent, as shown by the magenta portions of the bars. The total present value of path-independent counterparts for the payment distributions in the 30 years would be 96.23% as large as that of the actual payment distributions.

Figure 18
Cuatro: Present Values of Payments, Years 1 – 30

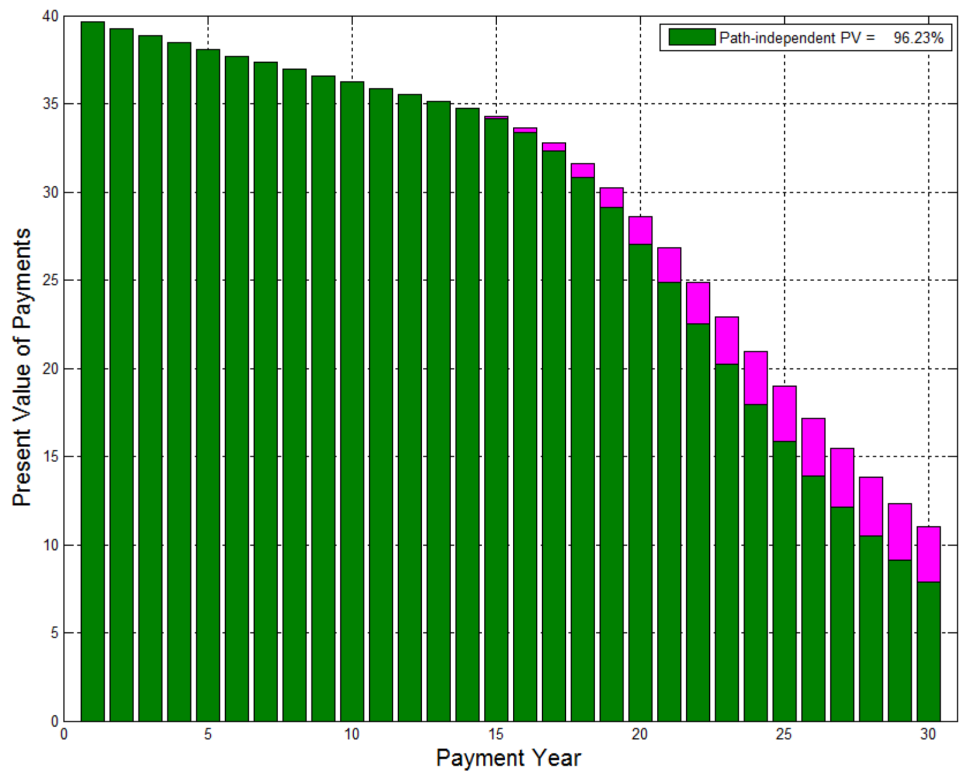

While in some cases payments will be smaller than desired or zero in the later years, in most of the possible scenarios, the four percent rule manages to provide completely smooth payments for all of the years in the stated horizon. But it does so, in large part, by usually spending less than the investor's total savings. This can be seen in Figure 19, which repeats the information in Figure 18 and adds the present value of the distribution of ending wealth (presumably unused by the investor). As can be seen, this is more than 10% of initial wealth. Moreover, it is highly path-dependent. Using path-independent counterparts to the distributions of the payments and ending wealth could reduce overall cost to slightly over 94% of the original amount.

Figure 19
Cuatro: Present Values of Payments, Years 1 – 30 and Ending Wealth

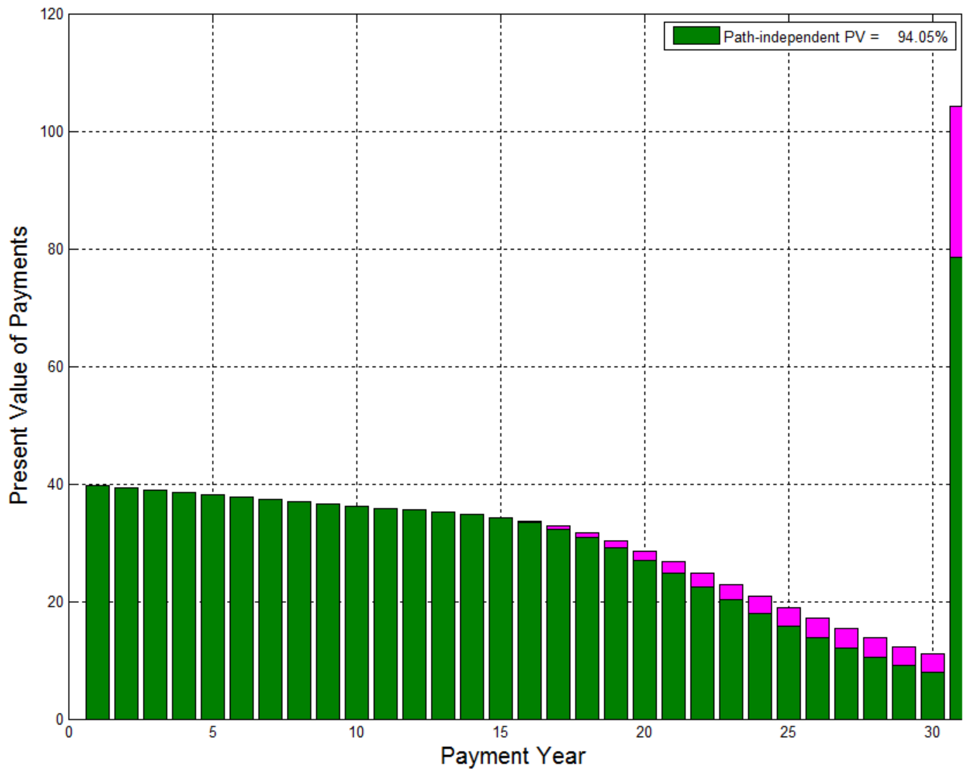

As argued in [Scott, Sharpe and Watson, 2009], the price of the four percent rule may be too high for many investors. Yet it remains popular among financial advisors. A cynic might suggest that some of them who recommend the approach may be mindful of the fact that it can provide substantial assets to be managed (and attendant fees) for many years.

BuyHo

The final example is drawn from a class of strategies that are simple, straightforward and easily understood by both advisors and investors. It uses lockboxes explicitly and minimizes transactions required to conform to the strategy chosen for each one. I call this version BuyHo to signify that it involves buying a combination of TIPS and the Market portfolio for each lockbox, then holding the initial positions up to the maturity date for the box. Of course as a practical matter, some transactions will probably be required to reinvest dividends and coupons in the market portfolio and to reinvest coupon payments from the TIPS portfolio. However, a true buy-and-hold alternative might be feasible if for each lockbox one could (a) buy a zero-coupon TIPS contract of the appropriate maturity and (b) invest in a market index fund that provides automatic reinvestment of all dividends. In the simplified capital market model used here, these issues are not relevant and the strategy is truly buy and hold.

Figure 20 shows the initial values of the two assets in each lockbox for the strategy. The darker bars represent TIPS and the lighter ones the market portfolio. The amounts were chosen using a very simple procedure. First, the total values for each lockbox were selected so that if each had been invested solely in TIPS, the payments would be identical for each of the thirty years. Next, the half the amount in the first box was invested in TIPS and the other half in the market portfolio. The proportion of the amount invested in the market for each of the subsequent boxes was then set to $0.50/\sqrt{t}$, where t represents the year for which the box is used to provide a payment.

Figure 20
BuyHo: Initial Values of TIPS and the Market Portfolio

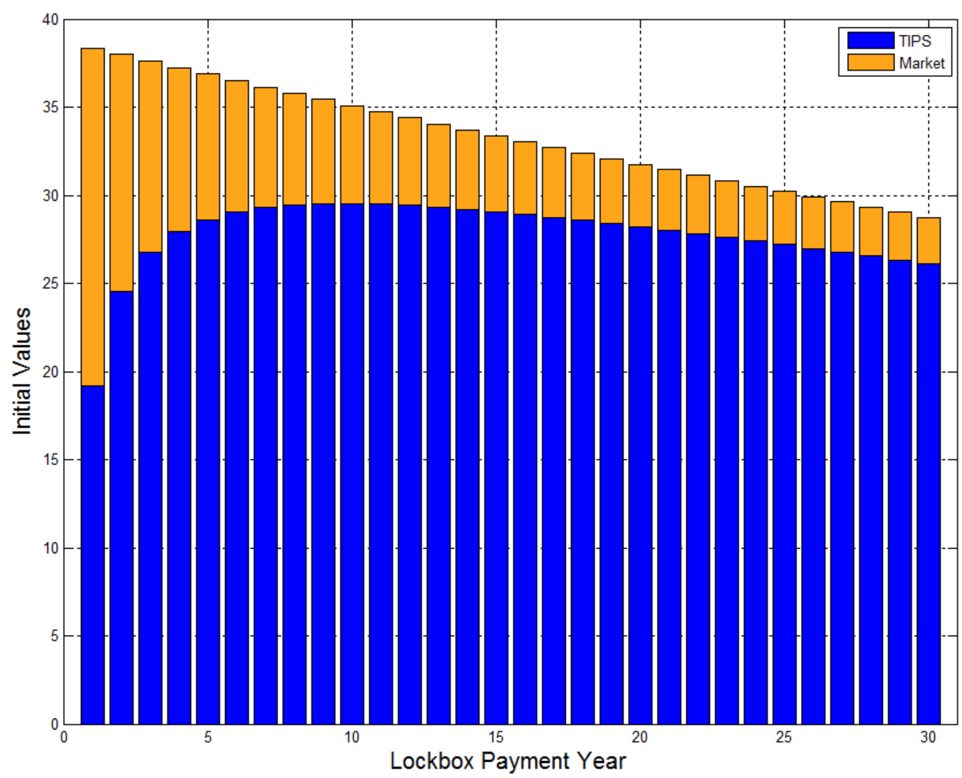

Figure 21 shows the resulting payment distributions. While the range is greater the farther in the future the year of payment, the rate of increase is substantially less than that for the Fido and Vido strategies.

Figure 21
BuyHo: Payment Distributions

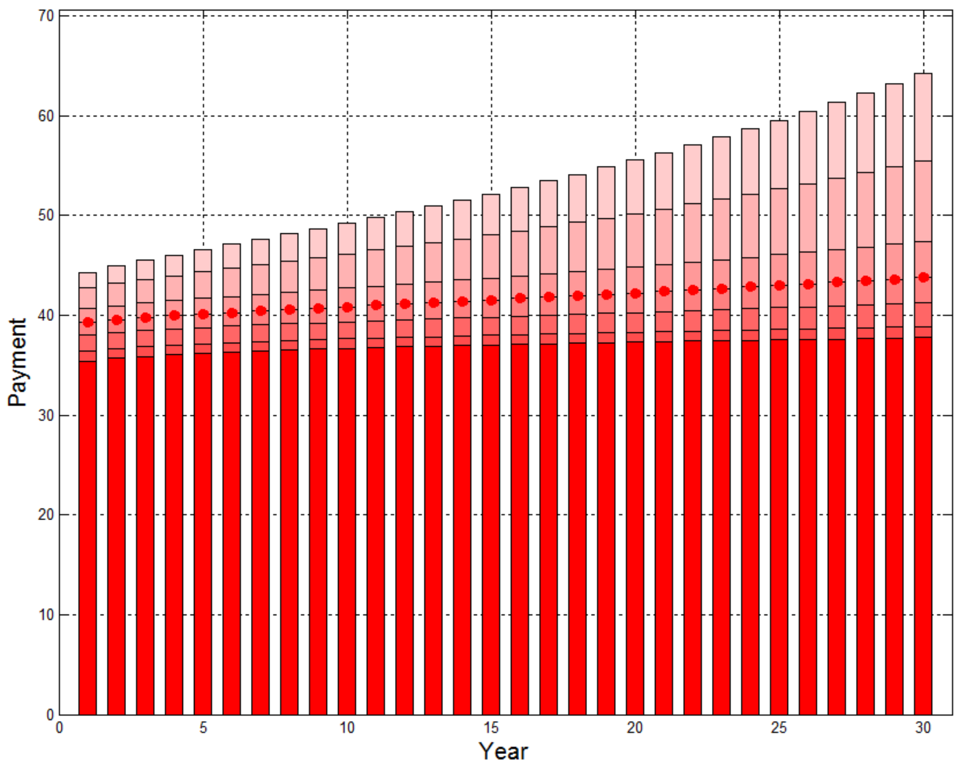

The payment ratio distributions, shown in Figure 22, also reflect the lower overall risk of the strategy. Changes from year to year are likely to be less wrenching than those of the Fido strategy. However, risk reduction comes with a cost -- the median payments are also less.

Figure 22
BuyHo: Payment Ratio Distributions

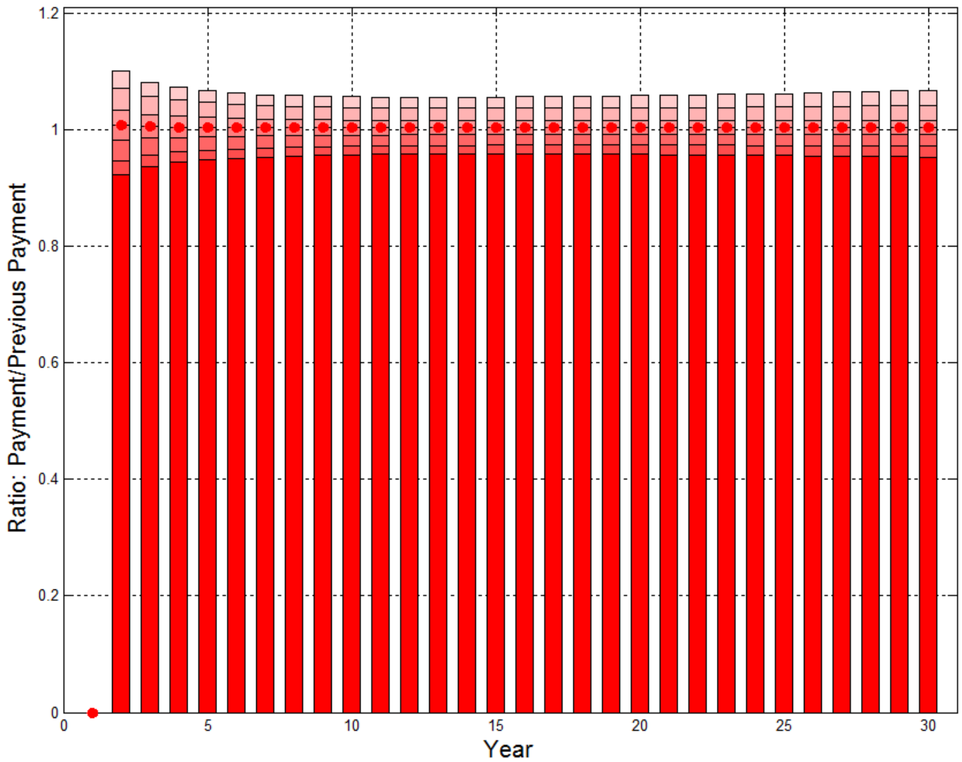

Buy and Hold strategies of this type lend themselves well to experimentation. One could build an interactive tool that would allow an investor to experiment with different allocations across lockboxes and, within boxes, with different allocations between risky and riskless assets (subject, of course, to the constraint that the sum of the initial values equal the investable wealth). By seeing the implications for payment distributions of different combinations of lockbox values and asset allocations within lockboxes, an investor might better understand the trade-offs involved and find it easier to select a preferred combination of possible outcomes.

Annuities

As indicated earlier, any of the strategies discussed and many more of the same genre could be used as a basis for an annuity contract. An insurance company could accept investments from a group of investors of a given age with similar mortality projections, charge each an initial lump sum, then promise to pay those still alive the payments dictated by the specified investment and spending strategy. Given a set of predicted mortality rates, the insurance company could estimate the actual value of the payments to be made in each year and invest the funds to provide those amounts. As a practical matter, of course, annuity providers charge higher fees than those implied by such calculations both to cover business expenses and to provide a reserve for adverse differences between actual and projected mortality rates.

To illustrate, Figure 23 provides a simple set of calculations based on a strategy that pays a constant real amount for each of the next 30 years. For each year the left-hand bar shows the present value of the payments if each were to be made without condition. The remaining calculations assume that an investor is male, 65 years old, obtains no utility from the prospect of leaving money to his children and/or charities, and thus chooses to purchase annuities from a company that charges 15% more than its expected cost based on current mortality tables. Clearly, it does not make sense for such an investor to pay the added costs for payments in years in which his chance of dying is small. Thus, for the first few years the right-hand bars in Figure 23 show the costs of investing directly, for the remaining years they show the costs of obtaining the payments from the insurance company. The overall total cost of this strategy is 68% of that of the non-annuitized alternative.

Figure 23
A Riskless Constant Payment Strategy:
Fido: Present Values with and without Annuitization

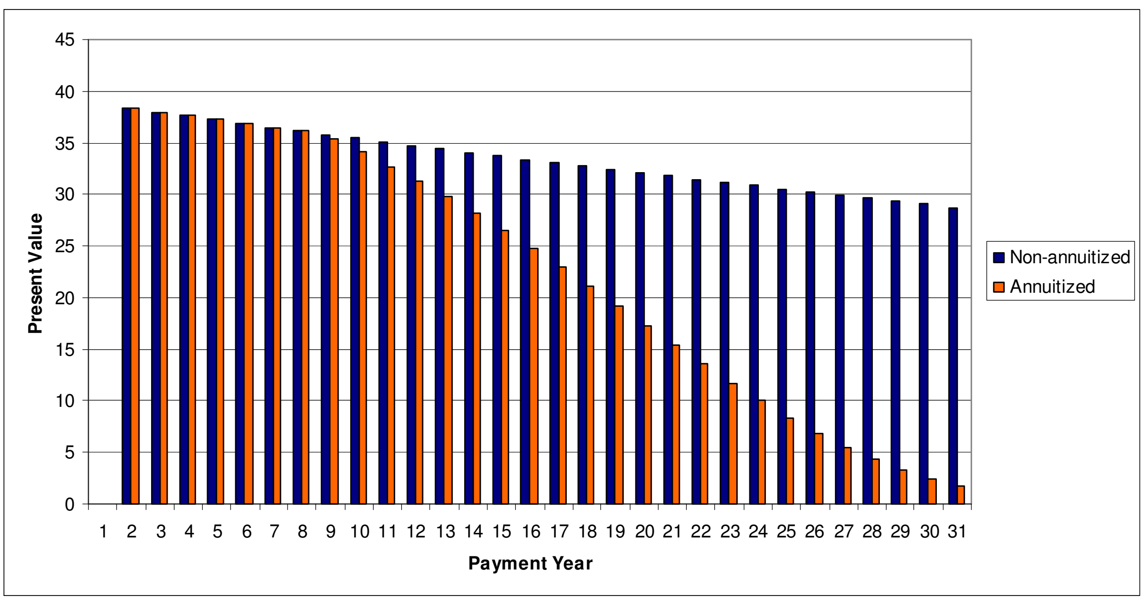

Some people may derive as much utility from contemplating spending money in future years on themselves as from thinking about the same amount of money being spent by their beneficiaries. Others may care only about money spent while they are alive. The former should avoid annuities and the latter should embrace them, at least for covering their later years. But many investors fall between these two extremes. For them, partial annuitization may be the best approach. Useful procedures for determining the best combination are beyond the scope of this paper. Suffice it to say that when selecting an overall approach it will be important to evaluate alternative investment and spending strategies for both the annuitized and the non-annuitized portions and that forecasts and valuations of the type presented here could help investors make informed choices.

Summary

There is more to be said about post-retirement financial strategies and much more work to be done. The variety of possible approaches offered by the financial community is already great and will undoubtedly increase as the number of retirees with significant amounts of wealth grows. It is important that such investors choose strategies that best meet their individual needs and desires. My goal in this paper has been simply to show ways in which a discrete-time, discrete-state simulation approach, combined with a pricing kernel valuation model might play a role in this process. Hopefully others will find these ideas worthy of extension and practical application.

References

Arrow, Kenneth J., 1952, "Le Role de valeurs boursiers pour la repartition le meillure des risques," Econometrie, Colloques Internationaux du Centre National de la Recherche Scientifique 11, pp. 41-47.

Bengen, William P., 1994, "Determining Withdrawal Rates Using Historical Data," Journal of Financial Planning, 7 (4), pp. 171-80.

Debreu, Gerard, 1959, *The Theory of Value*, Wiley and Sons, New York.

Dybvig, Phillip H., 1988a, "Distributional Aspects of Portfolio Choice," Journal of Business, 61, pp. 369-393.

Dybvig, Phillip H., 1988b, "Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market," *Review of Financial Studies*, 1 pp. 67-88.

Fidelity Income Replacement Funds[SM] 2011a, http://personal.fidelity.com/products/funds/content/WhatYouCanBuy/income_replacement_funds_overview.shtml.cvsr, April 9, 2011

Fidelity Income Replacement Funds Prospectus 2011b, Fidelity Investments, 2011.

John L. Maginn, Donald L. Tuttle, Jerald E. Pinto, Dennis W. McLeavey, ed. 2007, **Managing Investment Portfolios, A Dynamic Process**, 2007, CFA Institute

Merton, Robert C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," Journal of Economic Theory, 3, pp. 373-413.

NACUBO-Common Fund Study of Endowments, 2010. http://www.nacubo.org/Research/NACUBO_Endowment_Study.html

Scott, Jason S. with William F. Sharpe and John G. Watson, 2009, "The 4% Rule -- At What Price?", *Journal of Investment Management*, Vol. 7, No. 3, Third Quarter 2009, pp. 31-48

Sharpe, William F., 2006, "Retirement Financial Plannning: A State/Preference Approach," www.wsharpe.com/retecon/rfp.pdf, Feb. 2006.

Sharpe, William F., 2007a, **Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice,** Princeton University Press, 2007.

Sharpe, William F., 2007b, "Lockbox Separation," www.wsharpe.com/wp/lockbox.pdf, June 2007.

Vanguard Managed Payout Funds Prospectus, April 28, 2010.
http://www.interactivebrokers.com/prospectus/92205M101.pdf

Vanguard Withdrawal in Retirement Tool, 2011,
https://personal.vanguard.com/us/insights/retirement/withdrawal-in-retirement_tool

Von Neumann, John and Oskar Morgenstern, 1944, *Theory of Games and Economic Behaviour,* Princeton University Press, Second Edition, 1944.