

Retirement Income Analysis

with scenario matrices

William F. Sharpe

8. Valuation

What is the present value of a contract that promises to pay you \$1000 ten years from now? The answer depends on the conditions in which the promise will be kept. If the payment is absolutely certain to be paid it should be worth the current cost of a ten-year zero coupon bond paying \$1,000 at the time. But if there is any chance that the bond will pay less than \$1,000 or possibly nothing at all, it is worth less than the value of a risk-free bond. But how much less? What is it really worth?

This chapter provides a method for valuing strategies with uncertain future payments. Not surprisingly, many sources of retirement income fall into this category, especially when income is measured in real terms. As we will see, the ability to estimate the present value of a range of possible future real incomes can reveal important aspects of alternative retirement income approaches.

The Capital Asset Pricing Model

A standard approach for the valuation of uncertain future payments, taught in undergraduate and MBA level finance courses, is the Capital Asset Pricing Model described in Chapter 7. The key result for valuation is the Security Market Line, which implies that the expected return on an asset i should be determined by (1) the risk-free rate, (2) the scaled covariance of its returns with those of the market portfolio (its *beta* value) and (3) the excess expected return on the market over the risk-free rate. In symbols:

$$e_i = r_f + \beta_i (e_m - r_f)$$

In many cases, the expected returns and the risk-free rate are expressed as percentages (e.g. 10%) or fractions (e.g. 0.10). But they can also be expressed as value relatives (e.g. 1.10) and the equation will still hold. In what follows we assume that all returns are measured as value relatives.

Now, let v_i be the uncertain payment made by asset i at the end of a year. Its expected value relative will equal the expected value at year-end divided by the price today.

$$e_i = \frac{e(v_i)}{p}$$

For our purposes, the beta measure is based on the covariance of the value relative for the asset and that for the market portfolio:

$$\beta_i = \frac{cov(r_i, r_m)}{var(r_m)}$$

But the numerator can be expressed in terms of the current price of the asset and the covariance of the amount that will be received a year hence and the value relative for the market portfolio:

$$cov(r_i, r_m) = cov\left(\frac{v_i}{p}, r_m\right) = \left(\frac{1}{p}\right) \times cov(v_i, r_m)$$

Putting all these relationships together and re-arranging terms provides an expression for the current price of an asset based on the expected value of the probability distribution of the payments it will provide at year-end and the covariance of such payments with the value-relatives for the market portfolio:

$$p_i = \frac{e(v_i) - \frac{\text{cov}(v_i, r_m) \times (e_m - r_f)}{\text{var}(r_m)}}{r_f}$$

In the world of the CAPM this relationship will hold for any security or portfolio.

State Prices

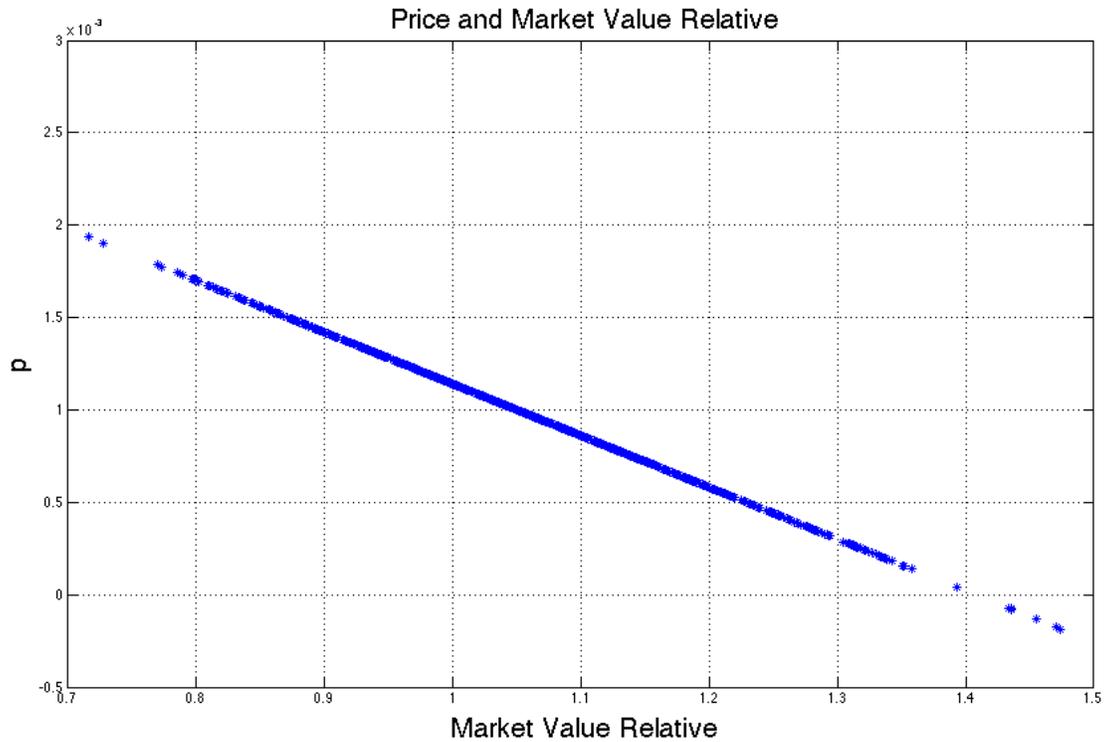
As indicated earlier, the CAPM is an equilibrium valuation model derived in the early 1960's based on the assumption that investors care only about the mean and variance of return distributions. And this assumption followed the prescriptions of Markowitz' portfolio theory, first published in 1952. A different approach to valuation of assets with uncertain returns was developed in the 1950's by Kenneth Arrow ("*Le Role de valeurs boursieres pour la repartition le meillure des risques*" in 1951) and Gerard Debreu (in a 1951 article and a 1959 book "*The Theory of Value: An Axiomatic Analysis of Economic Equilibrium*").

The Arrow/Debreu approach views the future in terms of a set of alternative *states of the world*. Their key insight was to show that in a *complete* market, for each future time period there could be a set of *contingent claims*, each of which would provide payment in one and only one of such states. The value of any security that provides payments that differ across states could then be determined by multiplying the amount paid in each state times the price of a claim to receive \$1 if and only that state occurs, then summing the results.

I explored the Arrow/Debreu approach at great length in my 2007 book, "*Investors and Markets, Portfolio Choices, Asset Prices and Investment Advice*". There I argued that the Arrow/Debreu (or *state-preference*) approach provides a richer way view asset valuation under uncertainty than does the CAPM . That said, the two approaches have considerable similarities. In the standard one-period setting for which the Markowitz approach was developed, and in which the CAPM is set, the mean/variance assumption can be considered a special case of the more general state-preference analysis. But it has some disadvantages in a one-period case and even more in a multi-period setting, as we will see.

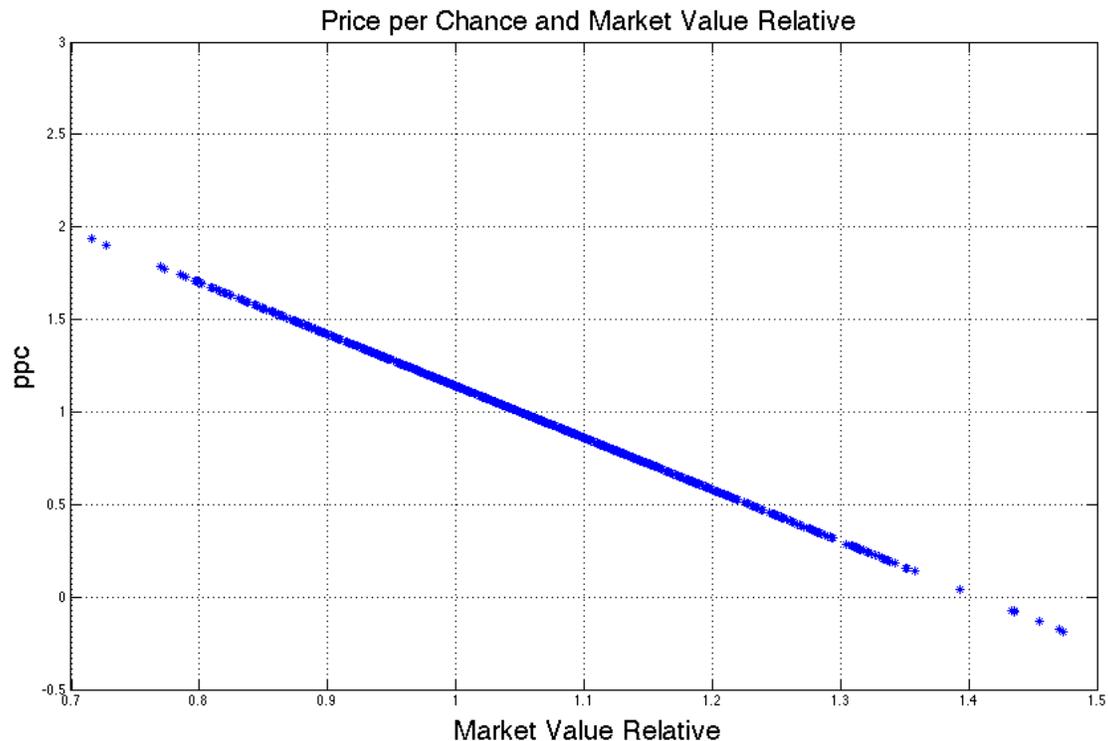
To start, consider a simple world with only one future period a year hence. For simplicity, assume there are 1,000 future states of the world and that we know the return on the market portfolio for each one. Now, consider an "Arrow/Debreu" security that pays \$1 at the end of the year if and only if the first scenario takes place and \$0 if any other scenario occurs. We can create its "payoff vector" with a "1" in the first row and "0" in the other 999 rows. Question: what is the *present value* of this security? Not a problem if the CAPM holds. We simply do the computations using the previous formula. The result is the *present value* of \$1 to be received if and only if state (scenario) 1 occurs. In finance parlance, it is the value of an *option* to receive \$1 if that condition obtains.

Now, assume that this procedure has been repeated for each of the 1,000 possible one-scenario payment options. The figure below shows the results from a case with our default market parameters (in which the risk-free value relative is 1.01, the excess return value relative for the market is 1.0425 and the standard deviation of the market value relative is 0.125).



Note, first that each of the *state prices* is small, since there is only one chance out of a thousand that any particular one will pay off. The y-value at top of the graph is 3×10^{-3} or \$0.003.

To get a better sense of the scale, it is useful to divide each price by the probability that the security will pay off (in this case, 1 out of 1,000) to obtain the *price per chance*, or PPC. The results for our example are shown in graph below.



Note that for every scenario with a given market value relative, the PPC value is the same. Moreover, the smaller the return on the market portfolio the greater is the PPC. More generally, the relationship is monotonic, downward-sloping and linear. The first two characteristics make great sense, as we will argue later. But the latter can lead to problems.

For states of the world in which the return on the market portfolio is especially high (here, greater than 40%) the state price and price per chance are both negative. This makes no economic sense. Why would someone actually pay you to hold an option which could either produce nothing (if the market return is less than 40%) or something (if it is greater than 40%)? A market in which you can obtain money now in return for accepting the chance (however small) that you will receive more money in the future is one that we can only dream about.

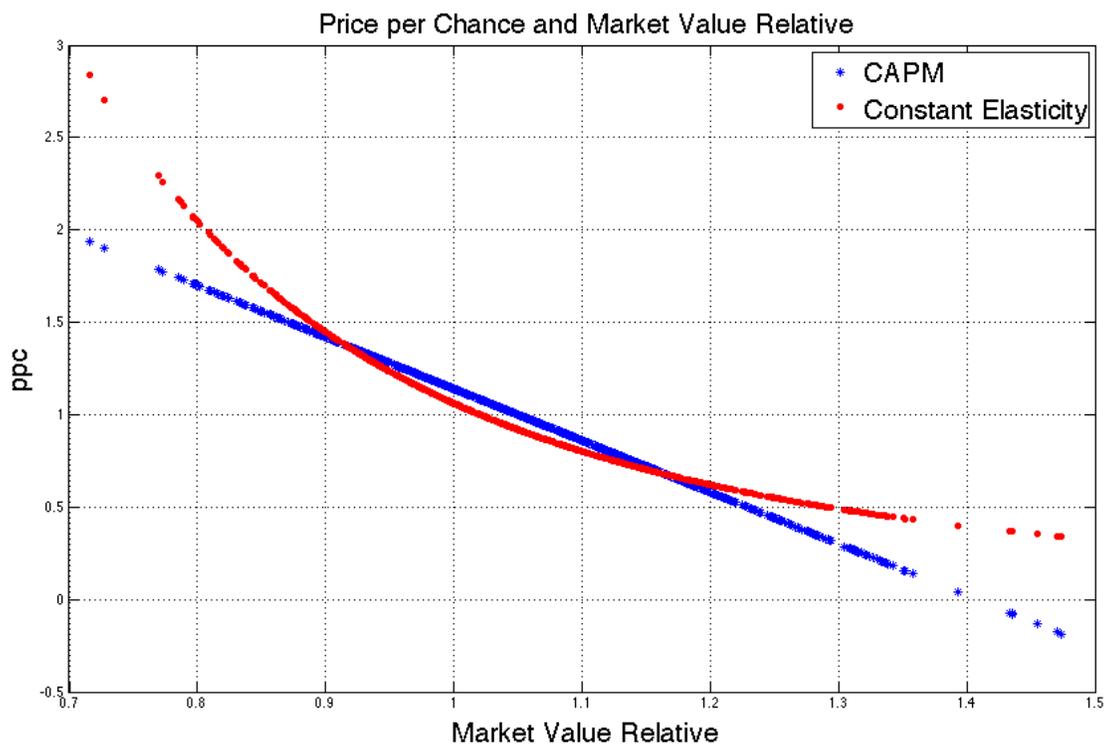
This problem comes from the assumption that investors care only about the mean and variance of the probability distribution of outcomes. We will have more to say about this in the next chapter. But such implications are well known. Markowitz himself has justified mean/variance preferences as only approximations for investors' true utility functions. Prices such these, obtained from the CAPM, are best viewed as approximations as well. For many purposes they may suffice. But, as we will see, a somewhat different approach is likely to provide better estimates of state prices and PPCs.

Constant Elasticity Pricing Kernels

In the asset pricing literature, the set of state prices (or prices per chance) is termed the *pricing kernel* for a market. Some call this (or the values obtained by multiplying each price by the risk-free value relative) the set of *Stochastic Discount Factors* (SDFs) or, collectively, the *Stochastic Discount Function*. We will avoid such terms, since they are at the very least confusing and could be misleading. Henceforth, the price for income to be received if and only if a state occurs will be termed the *state price* and the price per unit of chance (probability) the *PPC*.

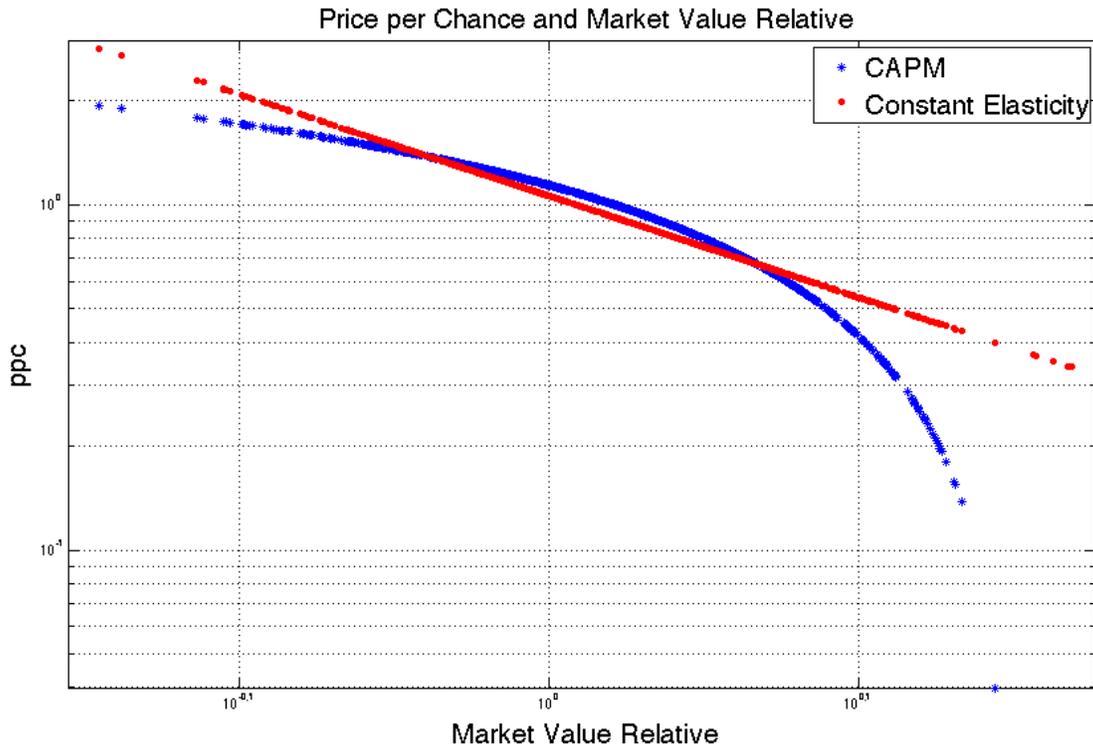
As we have seen, if the CAPM holds, the pricing kernel is linear. But this is at best a rough approximation. Negative prices make no economic sense. And one would imagine that the right to receive income in very dire markets (with low value relatives) should be worth considerably more than implied by the linear functions in the previous section.

The following figure shows an alternative (in red) along with the CAPM results (in blue).



As can be seen, this new pricing kernel produces relatively similar state prices for midrange market returns, but avoids negative state prices and also produces high state prices for extreme market declines – all features likely to be found in actual capital markets.

But how did we produce this new kernel? The answer is obvious when both axes are plotted on logarithmic scales (using the MATLAB *loglog* plotting function), as in the following diagram.



Here it is clear that the relationship shown by the red curve is linear, so that at every point in the diagram, a given small percentage change in the logarithm of the market value relative is associated with the same percentage change in the PPC. With a logarithmic scale, a given distance represents the same percentage change at every point (as can be easily seen by examining the horizontal grid lines). Economists use the term *elasticity* to refer to the ratio of the percentage change in one of two related variables divided by the percentage change in the other. In this case the (instantaneous) elasticity is the same at every point along the red curve. Thus the relationship exhibits *constant elasticity*.

In this case, the elasticity is -2.94, so that for every 1% increase in the market value relative the PPCs (and state prices) fall by roughly 2.94%. Later we will show how to calculate this coefficient directly. But first we focus on the economics of the situation.

Downward-sloping Demand Curves

Both our the linear pricing kernel and the constant elasticity version plot as downward-sloping functions, with lower prices associated with greater market value relatives. This makes great economic sense. In most markets, higher prices result in lower quantities demanded. Looked at the other way, scarce goods and services command higher prices in order to ration the existing supply. This is often termed the economic *law of demand*: in a diagram with quantity demanded on the horizontal axis and price on the vertical axis, the relationship will plot as a *downward-sloping curve*. Some would say this is the most important theorem in micro-economics. With rare exceptions, in competitive markets in which prices are freely set, lower prices are associated with greater quantities.

It seems entirely plausible that such a relationship should hold in capital markets. Low market value relatives are associated with *bad times* for investors and, in most cases for non-investors as well, since major declines in security values generally signal greater chances of hard times for the real economy. The pricing kernel should thus be downward-sloping for a very good reason: people will pay more for a scarce good (a dollar in bad times) than for a plentiful good (a dollar in good times). In bad times, there will be fewer dollars to go around, so people will pay more in advance to have one of them. This is the essence of asset pricing theory, whether it be the CAPM or this more general pricing kernel approach.

There is, of course, the question of how to estimate the actual demand curve. As we will see, there are good reasons for selecting the constant-elasticity form and, given our assumptions about the risk-free return and the distribution of market returns, the parameters of the function can be easily determined.

Multi-period Pricing Kernels

The setting for the CAPM and the mean/variance portfolio theory from which it was derived involves present investment followed by a payoff one period hence. More succinctly, both are one-period models. But in the real world, people often invest money now in order to receive payments not only a year from now but in subsequent years as well. To be sure, some investors have a horizon of one year (or less). But others have horizons of two, three or many years.

There is no agreed-upon model of equilibrium in a world in which investors have different horizons. But it is easy to show that the CAPM is not up to the task. Consider a world in which the model holds for both this year and next year. If so, the pricing kernel for each year will be a linear function of the value-relative for the market portfolio in that year:

$$p_1 = a - b R_{m1}$$
$$p_2 = a - b R_{m2}$$

The price (present value) of \$1 two years from now will be the product of its present value for year two times the present value for year 1:

$$p_1 p_2 = a^2 - abR_{m1} - ab R_{m2} + b^2 R_{m1} R_{m2}$$

Clearly the price today of a dollar two years from now will depend not only on the total value relative for the market over the two years (the product of the two value relatives in the last term) but also on the way in which that total value was achieved (the individual returns in the second and third terms on the left of the equal sign). And the farther in the future the payment, the more terms that will be required for its valuation.

Contrast this with the case when the pricing kernel has constant elasticity. Assume that the one-period kernel is:

$$p = a R_m^{-b}$$

Then for a horizon of two years:

$$p_1 p_2 = a R_{m1}^{-b} \times a R_{m2}^{-b} = a^2 (R_{m1} R_{m2})^{-b}$$

More generally:

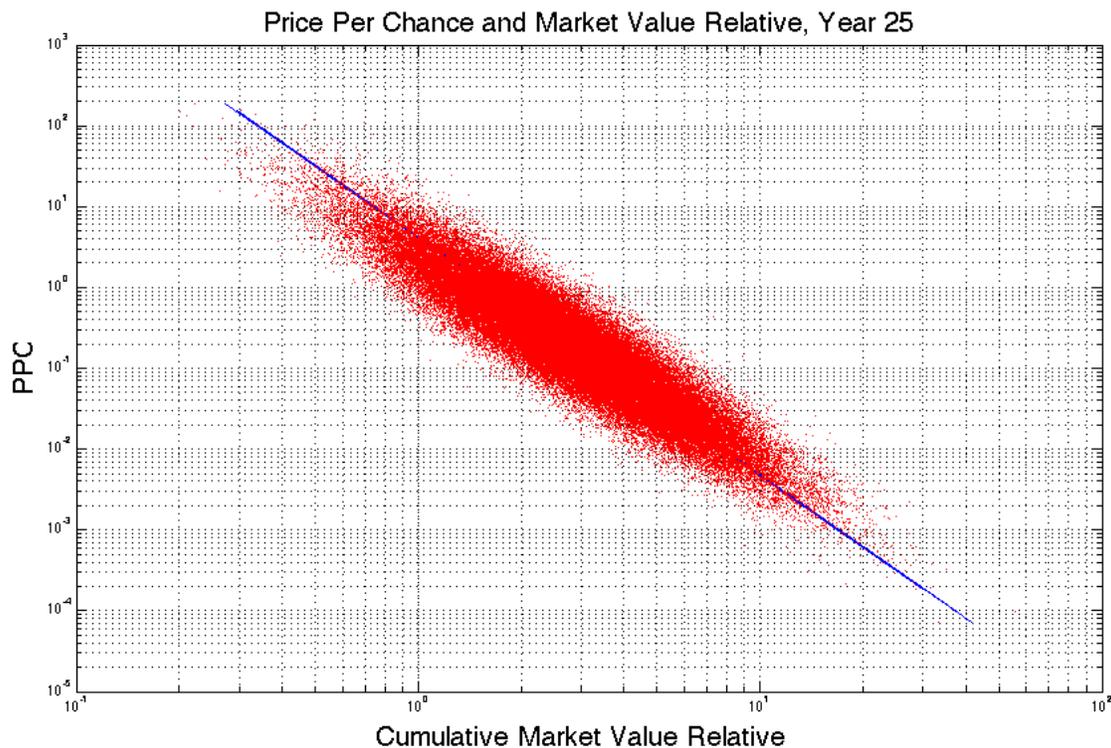
$$p_{st} = a^t R_{mst}^{-b}$$

where p_{st} is the present value today of a dollar to be received at time t in scenario s and R_{mst} is the cumulative value relative for the market from the present to time t in scenario s .

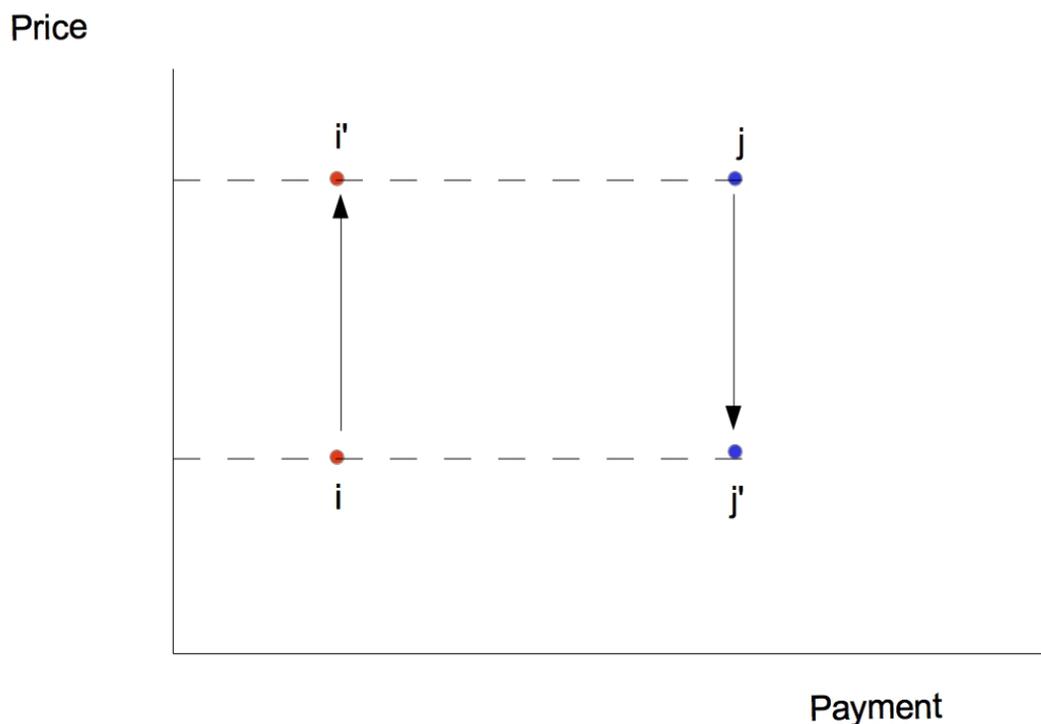
Graphically, the relationship between state price and cumulative market return will be a downward-sloping curve for any future horizon (t). This has important implications, as we will see.

Cost-efficiency

The figure below, based on 100,000 scenarios and returns for 25 years, shows PPC values and Cumulative Value Relatives for the final year, using a constant-elasticity pricing kernel based on our standard parameters for the risk-free return and the distribution of market returns. As before, the values are plotted on logarithmic scales. Two strategies are represented. The first, which is a “buy and hold” strategy that invests in the market portfolio at the outset and holds it until year 25, plots as a set of points on the straight blue line. The second, represented by the red points, involves an “active management” strategy in which some assets are held in less-than-market proportions and others are held in more-than-market proportions. This could be done on a permanent basis, say by holding only one of the four components in our world bond/stock market surrogate. Or a manager might choose different holdings in each year and scenario. Or a combination of the two approaches. For this exercise, returns for the active strategy were obtained by adding to each market value relative in the matrix a normally-distributed variable with a mean of zero and a standard deviation of 0.05.



Clearly, such an active strategy has non-market risk, since the red dots scatter around the pricing kernel. And, in an important sense, this is an *inefficient* strategy. Consider any case in which one of two points lies to the northeast of the other. The figure below provides an exaggerated example.



For emphasis, the horizontal axis has been labeled “payment” since the value relative can be used as a payment and the vertical axis has been labelled “price”, since a PPC is simply a present value (price) divided by its probability. Here scenario j plots to the northeast of scenario i , showing that it provides a greater payment in a state with a higher price. Consider a switch in which payment i is provided in the scenario with a greater price and scenario j is provided in the scenario with a lower price. The resulting situation, shown by points i' and j' , provides the same two payments but the total cost is clearly lower.

Such an active strategy is clearly inefficient in this sense, since many pairs of points can be found in which the greater of the two payments is provided in the more expensive state of the world. In any such case there is a better way to provide the same set of payments at lower cost. All one has to do is sort the prices from lowest to highest, then arrange to get the highest payment in the least expensive state, the next-to-highest payment in the next-to-least expensive state, and so on. More generally, one can sort the vector of prices in ascending order and the vector of payments in descending order, then assign each element in one vector to the one in the same position in the other. The sum of their products will be the cheapest way to obtain the original set of payments.

More simply, the following Matlab code will do the job.

```
currentValue = prices' * payments;  
minimumValue = sort( prices, 'ascend' )' * sort( payments, 'descend' );
```

Finally, we can divide the minimum value by the current value to provide a measure of the *cost efficiency* of a strategy:

```
costEfficiency = minimumValue / currentValue;
```

In this case shown earlier, the the cost efficiency of the active strategy shown by the red dots is 0.9185, indicating that the same exact distribution of payments could have been obtained for 91.85% of the cost of the current strategy. To cover the possibility of ties, we can say that:

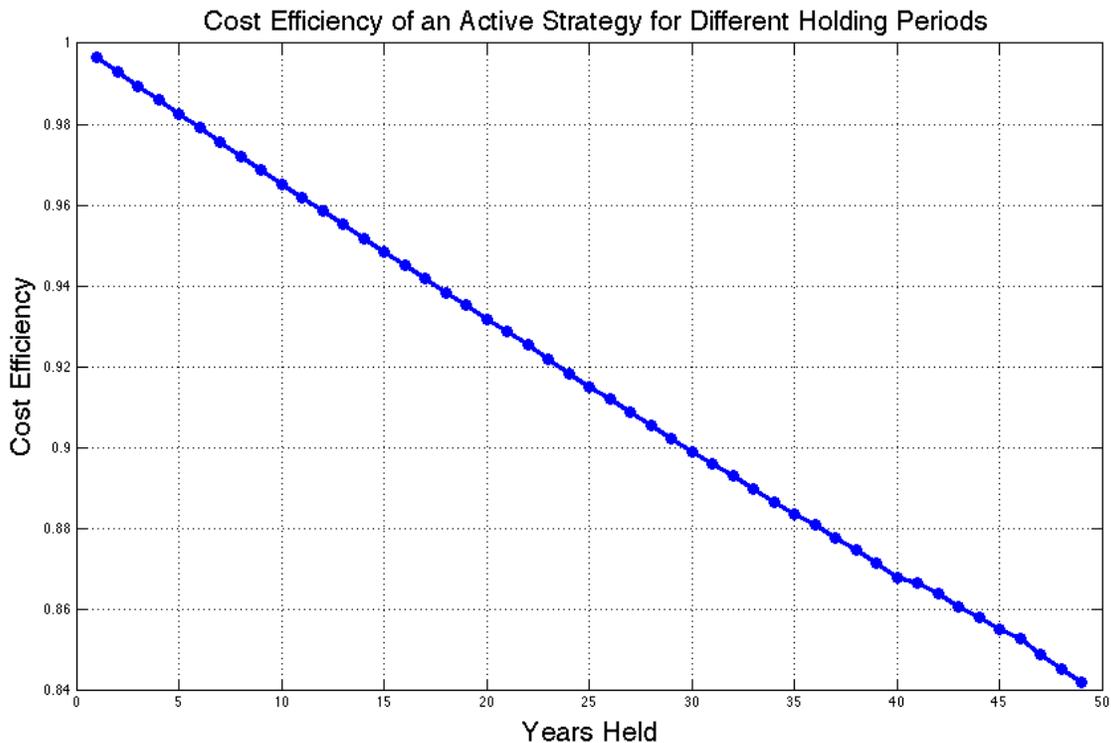
for a cost-efficient strategy, payments are a non-increasing function of state prices.

And since, market value relatives are a non-increasing function of state prices:

for a cost-efficient strategy, payments are a non-decreasing function of market value relatives.

For conciseness we term an approach of this type a *market-based strategy* – its payments do not have plot as a strictly upward sloping function of market returns or value relatives, but the curve can never go down.

If an active manager simply adds uncertainty to returns by overweighting some investments and underweighting others investments relative to market proportions, the investor will obtain results that could have been produced for less with a market-based strategy. In the earlier case in which the manager provided a return each year equal to that of the market return plus a normally-distributed variable with a mean of zero and a standard deviation of 0.05, the result could lower a recipient's standard of living 25 years hence by over 8%. And this is in addition to the losses resulting from management fees, transactions costs, etc.. Of course the impact of active management will depend on the period over which it does its damage. The following graph shows the cost-efficiency in this case for holding periods from 1 to 50 years in length. The longer the period, the greater the possible loss.



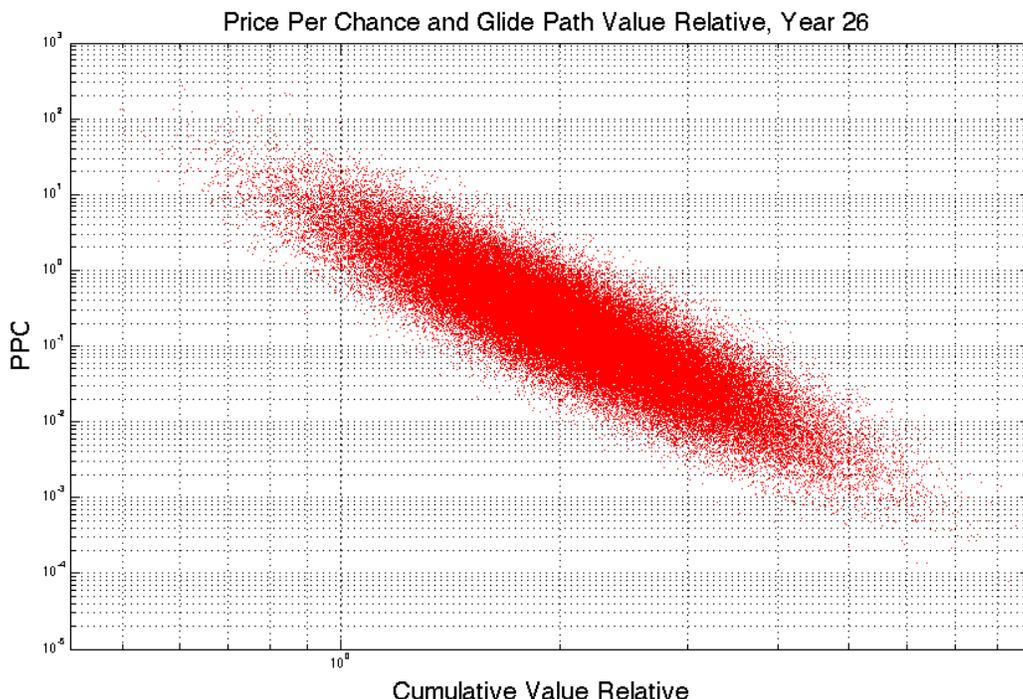
Of course these estimates are based on the assumption that our market portfolio is in fact “priced” in capital markets so that any other portfolios will be cost-inefficient. And the quantitative impact will depend not only on the length of time the portfolio is held but also on the magnitudes of departures from market returns, which could be less (or more) than assumed in our example.

Investment Glide Paths

One might assume that any strategy which invests only in the market portfolio and/or TIPS would have a cost-efficiency of 1.0 (100%). This will indeed be the case for any approach that invests in one or both assets, holds each component for some number of years, then spends the proceeds at the end of the period. But not necessarily for an approach that varies holdings in the two assets from time to time before the proceeds are spent .

Consider an investment policy that starts with 100% invested in the market portfolio, then buys and sells securities at the end of each year so that the proportion in the market portfolio will be 4% less than at the end of the prior year, ending with 0% invested in the market portfolio in year 25 and thereafter. *Glide path* strategies of this type are often recommend for those accumulating assets for retirement and sometimes for retirees as well.

Here are results for year 25 with 100,000 scenarios, using our standard return assumptions. Clearly, there is not one-to-one relationship between the amount of an ending value and its price. Why? Because there is not a monotonic relationship between the cumulative performance of the market portfolio and that of the strategy. The strategy's performance depends on both the terminal value of the market portfolio but also the path that it took to reach that value. More succinctly, this is a *path-dependent* strategy. And, as in our previous case, the strategy is not cost-efficient.



Making the computations shown in the previous section, yields an estimated cost efficiency for year 26 of 0.926. This means that it would be possible to achieve the same probability distribution of returns with a market-based strategy for 92.6% of the cost.

To be sure, these results depend on our overall model of capital market returns and valuations as well as the specific parameters that we have used for the computations. Different assumptions will undoubtedly give different numeric estimates. And shorter and more gentle glide paths will generally cause less inefficiency (as we will see in later analyses). But many financial advisors advocate strategies for providing retirement income that give path-dependent returns; thus it is useful to estimate the resultant inefficiencies and added costs.

A Multi-period Capital Market Equilibrium

Consider a world in which the pricing kernel for each period is the same and has constant elasticity. As we have shown, only investment strategies that provide payments in each year that are monotonic non-decreasing functions of the return on the market portfolio are cost-efficient. In such a setting, each informed investor will choose to invest in either the riskless asset, the market portfolio, a combination of the two or some security with returns that are market-based (in the sense that the returns are a non-decreasing function of the return on the market portfolio). Assume that there are investors with different horizons, from one year to possibly many years. In such a setting, every informed investor would choose some combination of a riskless asset and the market portfolio or a derivative thereof. At any given time, the markets would clear and the market portfolio would be cost-efficient for every horizon (t) since as we have shown:

$$p_{st} = a^t R_{mst}^{-b}$$

This is a powerful reason for choosing such a pricing kernel. It is consistent with a possible set of equilibrium asset prices.

Of course this convenient result doesn't mean that actual asset prices are set in this manner. In the real world, people do have different horizons and in many cases they tailor their portfolios accordingly in the belief that the results are preferable to the sorts of market-based strategies that we recommend. Actual aspects of market equilibrium are undoubtedly far more complex than those in our simple model thereof. Yes, our assumptions are consistent with equilibrium, but as Ralph Waldo Emerson, a nineteenth century American poet and philosopher wrote “a foolish consistency is the hobgoblin of little minds, adored by little statesmen and philosophers and divines.” He could have added to the list economists such as the present author.

One might imagine that empirical evidence could be marshaled to support or reject the characteristics of our assumed equilibrium. Some attempts have been made, but the problem is a difficult one. First, historic empirical data are limited. Even if underlying return distributions were constant from year to year, we have a limited number of observations of annual returns on a broad market portfolio (say, 100 from a 100-year history). And we have no direct way of measuring *ex ante* prices for market returns that did not occur.

An alternative is to observe prices of securities such as traded option contracts that offer payments related to the return on a broad market portfolio. But such prices typically reflect both state prices and probabilities. In our world of many scenarios, it is possible to determine a price for each state and time based on the market return at that time in that state. And each of our scenarios has the same probability. And there may be multiple states with the same market return at a given time. The price for a claim that pays one dollar when the market return equals that amount will reflect both the number of such states and the price per state. To be sure, it will have the same price per chance (PPC) as does each of the component states and one may be able to infer that PPC value from the prices of different option contracts. But to find the state price, one must make an assumption about the forecast probability of the state – the chance (C) in the denominator of price per chance (PPC).

Valiant efforts to overcome these obstacles have been made by several researchers using the prices of options on U.S. stocks and bonds – most notably, work by Stephen Ross in “The Recovery Theorem, in the *Journal of Finance* in 2015. But data are limited and the task is not a simple one. At this point results, are suggestive but not definitive. That said, the pricing kernels derived in such studies appear to exhibit decreasing slopes in a standard diagram with price on the vertical axis and payoff on the underlying asset on the horizontal axis, as does our constant-elasticity formula.

Computing the Pricing Kernel Parameters

Given the form of our pricing kernel, it remains only to compute the parameters associated with the assumptions about returns on the market portfolio and the riskless asset. We need to find the values of parameters a and b for the pricing kernel:

$$p_{st} = a^t R_{mst}^{-b}$$

These parameters should produce present values (prices) that correctly value the total proceeds in any future year from investing a dollar in the market portfolio at \$1. And the same present values should value the total proceeds in any future year from investing a dollar in the risk-free asset at \$1. Of course, if this holds for the first year it will hold for every future year and, as a result, for any multi-year holding period. Thus it suffices to find parameters that “price” returns for the first year:

$$p_{s1} = a^1 R_{m1}^{-b}$$

More simply put:

$$\sum p_s R_{ms} = 1$$

$$\sum p_s R_{fs} = 1$$

But since the two left-hand sides equal the same value (1), we may write:

$$\sum p_s R_{ms} = \sum p_s R_{fs}$$

Substituting our pricing formula gives:

$$\sum a R_{ms}^{1-b} = \sum a R_{ms}^{-b} R_{fs}$$

Clearly, coefficient a can be dropped from each side, leaving an equation with a single unknown – the elasticity measure b . It would be a relatively simple matter to solve for this value numerically, say by trying different values until the equation held to any desired degree of precision. Then the value of coefficient a could be found that would make each side in the original equation equal 1.

This procedure could be repeated for the each year. If there were an infinite number of scenarios, the resultant estimates of the coefficients should be the same. But even with 100,000 scenarios we do not have the entire population of possible returns on the market that would be generated by the assumed lognormal distribution. Instead, we have a samples of that population, one sample for each year. Because of this, the method will provide different values of the coefficients a and b in our pricing equation for each year in the analysis. And this, in turn, will violate the desired characteristics of our multi-period equilibrium.

Fortunately there is another approach, provided by one of my coauthors John Watson for an early paper employing this valuation method: “The 4% Rule – at What Price?” Jason S. Scott, William F. Sharpe and John G. Watson, in the *Journal of Investment Management*, Third Quarter, 2009. It computes the coefficients for a constant elasticity pricing kernel directly from the parameters of the lognormal return distribution for the market return and the risk-free rate.

The first step is to produce the constant-elasticity coefficient:

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

Given this, the constant term can be computed:

$$a = (\sqrt{E_m \times R_f})^{b-1}$$

Our standard assumptions are that:

$$\begin{aligned} R_f &= 1.01 \\ E_m &= 1.0525 \\ S_m &= 0.125 \end{aligned}$$

Which gives:

$$\begin{aligned} a &= 1.0524 \\ b &= 2.9428 \end{aligned}$$

Computing Present Values

All the results from the previous derivations in this chapter are implemented by adding six statements to the *market_process()* function.

Recall that in chapter 7 we included a statement to compute the total expected return for the market portfolio by adding the risk-free rate of return to the market expected excess return:

```
u = market.exRm + ( market.rf - 1 );
```

Using this parameter, the standard deviation of the market return and the risk-free rate, we can compute the parameters of the pricing function:

```
b = log( u/market.rf ) / log(1 + (market.sdRm^2) / (u^2) );  
a = sqrt( u*market.rf ) ^ ( b-1 );
```

Next, we create a set of a^t terms as a vector:

```
as = ones(nrows,1) * ( a .^ ( 0: ncols-1 ) );
```

then a complete matrix of price per chance values:

```
market.ppcsM = as .* ( market.cumRmsM .^ -b );
```

To complete the task, we divide by the number of scenarios to obtain a matrix of present values:

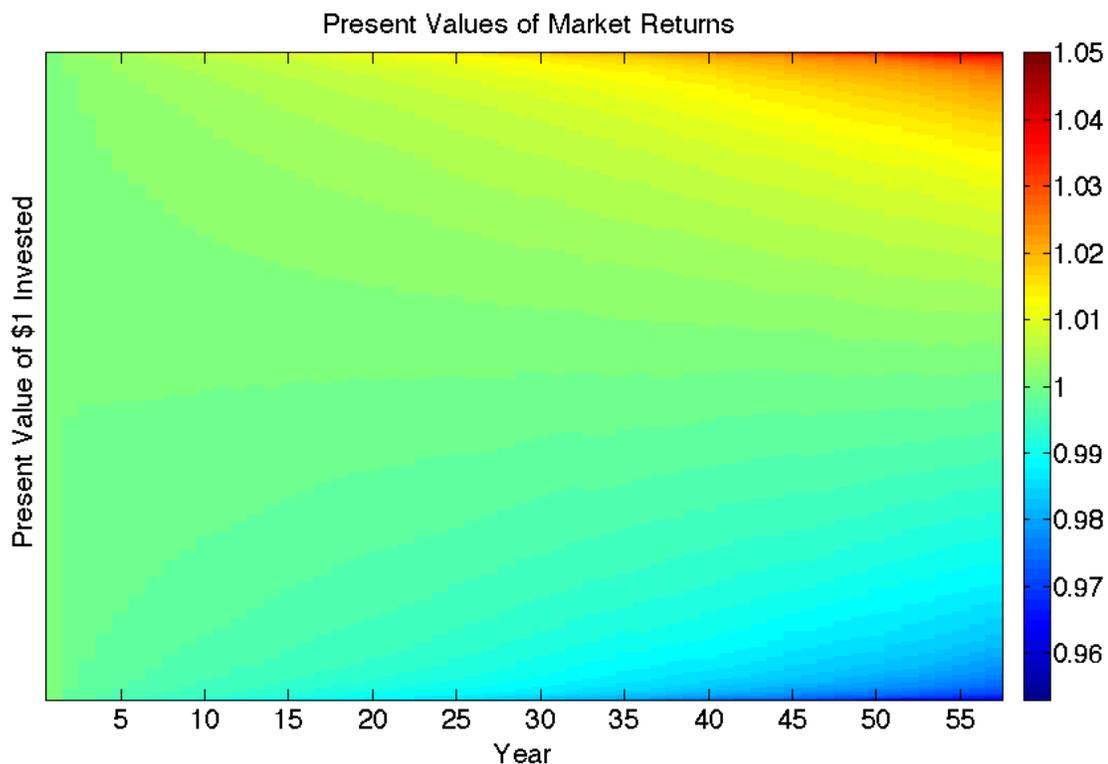
```
market.pvsM = market.ppcsM / nrows;
```

Sampling Errors

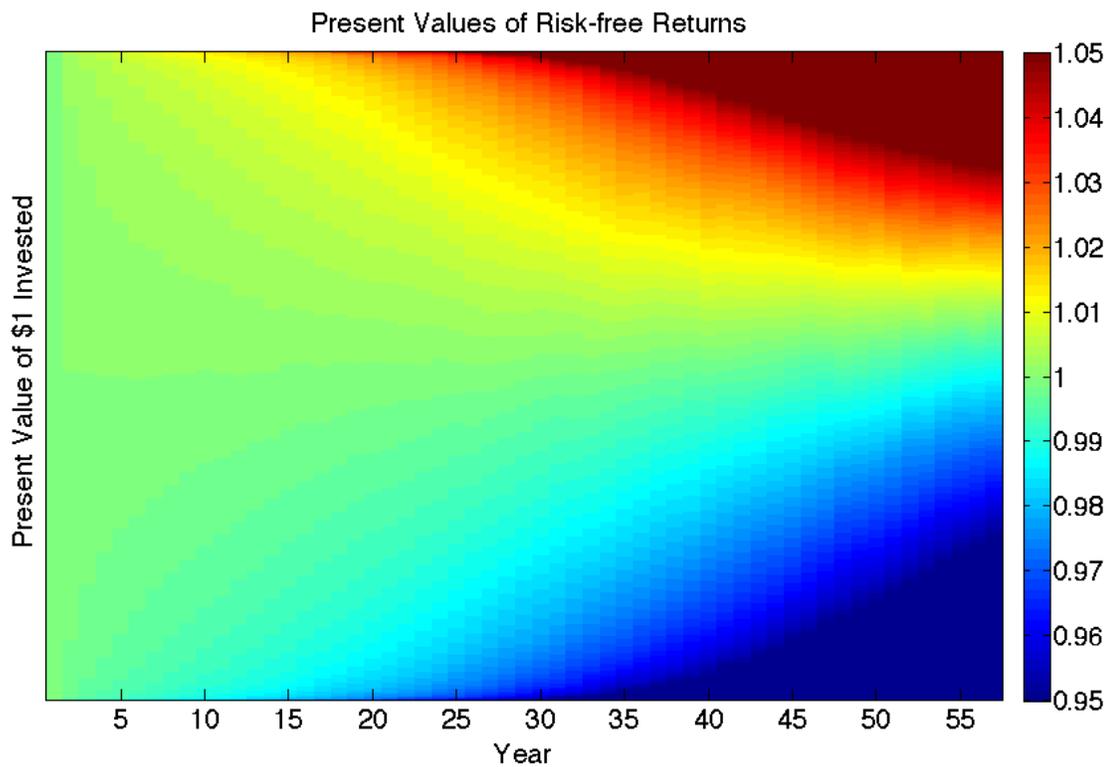
It might seem that the creation of 100,000 scenarios is overkill. One often hears about *Monte Carlo analyses* in which 1,000 or 5,000 or even 10,000 scenarios have been analyzed. We do not use the term “*Monte Carlo*” since it conjures up images of super-rich people gambling, not ordinary people trying to live comfortably in retirement. That said, many such studies deal with probability distributions of market portfolio returns, as do our analyses. And there is good reason to believe that results of such analyses would have been more valuable had more scenarios been generated. As we have seen, the time, effort and computer time involved need not be significantly great as long as programs are written in a language able to efficiently perform operations on large matrices (cue the advertisement for Matlab).

This said, even with large samples there will still be some differences between the characteristics of the samples drawn from underlying probability distributions and those of the hypothesized population parameters, as examination of the present values for our pricing kernel will show.

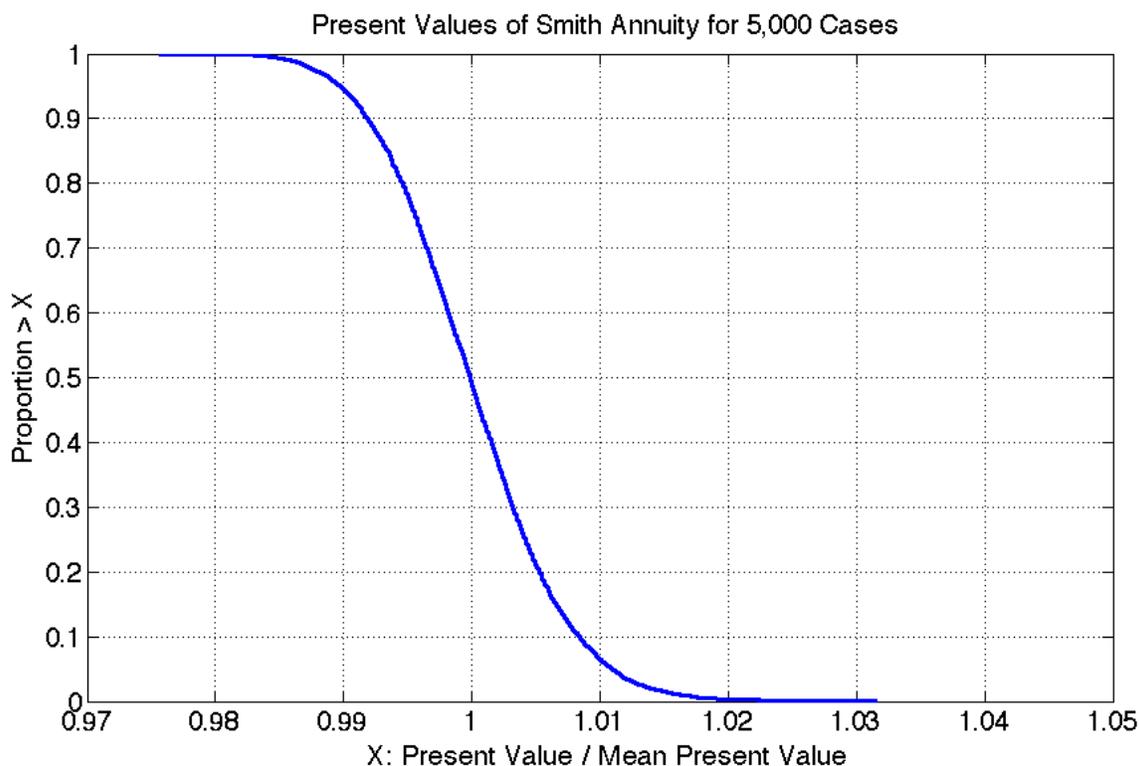
In principle, the present value of the cumulative returns on the market portfolio in any year should equal \$1.00, and so should the present value of the cumulative returns on the risk-free asset. But for any particular set of draws from the distributions this may not be the case. The following heat diagram shows the results from 5,000 analyses, each of which used 100,000 scenarios. For each future from year 1 through year 57 (the last for the Smith Case), the colors show the distribution of the computed present value for investment in the market portfolio. In the majority of cases it is close to \$1.00, and in almost all cases the values fall with a range from \$0.95 to \$1.05 (and the very few outside the upper and lower values are included with those values). But for the very distant years, the computed value can differ substantially from the theoretical population value of 1.0.



The next diagram provides results for the risk-free rate of interest in these cases. Although there is no uncertainty about the cumulative value of a risk-free investment, divergences of the market return distribution from its theoretical values produce errors in the pricing kernel. These errors may not be fully offset by the divergences in the market return distribution as in the previous diagram, and such price errors can lead to greater possible errors in the valuation of long-term risk-free investments. As in the previous diagram, the range of possible sample errors is greater for more distant years and can be substantial for cash flows to be received in the very distant future.



Of course, the probabilities of receiving substantial incomes after 50 years in retirement are relatively small, due to the inexorable toll taken by mortality. Thus the actual impact of sampling errors on valuations for retirement income plans is likely to be relatively small. To provide a more relevant view, we valued a fixed annuity for the Smiths that provides a constant real income for every year that Bob and /or Sue is/are alive plus an amount equal to one-year's payment to their estate. The figure below shows the results from 5,000 cases (each with 100,000 scenarios) of the present value divided by the mean over all the cases. As can be seen, the values fall around the theoretical (mean) value with relatively small deviations. In no case was the error more than 2%, and in roughly 90% of the cases it was less than 1%.



Another mitigating fact is that in many cases our interest is in the *relative* values of various aspects of a retirement plan (such as the present value of fees vis-a-vis spendable income) so sample pricing errors in some variables will be at least partially offset by similar errors in other variables.

The bottom line is that valuations will be subject sampling errors. In some cases it may be prudent to increase the sample size (say, from 100,000 scenarios to 500,000) or to repeat the analysis to examine the variation between cases. It may also be useful to set the random number generator so that results may be repeated in future analyses. The simplest way to do this is to execute the Matlab command:

rng(n)

where n (the “seed”) is a positive integer. This will produce a predictable set of “pseudo-random” numbers for each of the future uses of *rand* or *randn* functions. We will generally not do this, retaining a certain element of surprise whenever an analysis is re-run but accepting the fact that these are, after all, attempts to model phenomena that are ultimately less-than-perfectly known, even probabilistically.

With this caveat in mind, we are ready to move on to considerations of the preferences of the beneficiaries of a strategy or strategies for providing retirement income.