Retirement Income Analysis
with scenario matrices

William F. Sharpe

9. Utility

Assessing Retirement Income Scenarios

The goal of this book is to show how a matrix of scenarios for possible retirement income over a number years can be generated and assessed. Future chapters will discuss different strategies for producing retirement income and their properties. But ultimately, recipients need to choose a strategy or combination of strategies and the associated parameters thereof. The goal is, of course, to find the approach that will be "best" for those who will receive the income. To be (only) slightly more specific, we wish to find the feasible income scenario matrix that will maximize the recipients' "happiness". All (!) we need is a measure relating such happiness to the elements in an income scenario matrix.

Cognitive psychologists, including those in fields such as *behavioral economics* and *behavioral finance* have shown that human beings are not super-rational computing machines. Instead, they often make choices that are internally inconsistent, seemingly dominated by alternative approaches, and almost impossible to represent as the maximization of some well-formulated function.

Some argue that the only feasible way to help people choose among alternative scenarios for retirement income is to show them summary measures of each of two scenarios, ask them to pick the preferred one, repeat the experiment with the chosen scenario pitted against a third, then continue to pair the winner of each contest against a newcomer until it seems reasonable to stop. One can think of this as the *optometrist approach,* in which two lenses are tested at a time, with the chosen lens paired against another after each trial ("which is better, lens A or lens B? B? O.K., which is better, lens B or lens C", etc.). Of course, it isn't simple to asses a matrix with several million elements, let alone compare it with an other of equal size. We will develop some graphical portrayals that can help, but the task can still be arduous.

This chapter focuses on approaches advocated by traditional economists, who have assumed the existence of rational decision makers. Richard Thaler and Cass Sunstein in their book *Nudge: Improving Decisions About Health, Wealth and Happiness,* have differentiated between such mythical "Econs" and actual "Humans" (real people). While their points are very well taken, it is at least instructive to see whether traditional approaches can be at all helpful, while still recognizing that retirees are in fact human and need to be intimately involved in the choice of a retirement income strategy.

With these caveats in mind, we turn to the concept of *utility* and methods for maximizing it in the context of retirement income.

Utility

To take the broadest view, our present goal is to compute a numeric value for the “expected utility” of any retirement income scenario. The idea is that *utility* is some measure of “happiness” and *expected utility* is an average of all the possible levels of future happiness, weighted by their probabilities. The goal is thus to pick from feasible alternative scenarios the “best” one that has the greatest expected utility.

Throughout this book, we will assume that all income received at the beginning of a year is spent in that year. This allows us to regard utility as a function of income, which is equivalent to the usual economic formulation in which utility is a function of consumption. More generally, any plans to save income for a future year or to spend amounts previously saved should be included in the specifications of an overall retirement income strategy so that the amounts shown for income in a year equal the consumption in that year. We will assume this is the case and hence that the income from a strategy or combination of strategies will be the entire source of utility in that year.

Now, letting *yM* be a matrix of income with the dimensions of our previous matrices, the goal is to:

maximize: EU(yM)

subject to: C(yM) <= B

where:

EU(yM) = the expected utility of income matrix yM

C(yM) = the cost (present value) of the income payments in matrix yM

B = the recipients' budget (wealth) used to provide the income payments yM

The expected utility function *EU( )* will depend (at the least) on the income in each cell of the matrix, the column (year) in which it is received, and the personal state at the time in that scenario.

Behavioral economists generally scoff at this formulation, and with considerable justification. Nonetheless, some useful lessons can be learned by exploring its possible implications. We will do so in stages, beginning with a simple one-year setting in which all recipients survive to receive the available income.

Maximizing Utility in a One-Period Setting

The standard view of utility is that it increases as income increases, but at a decreasing rate, as in the following figure.

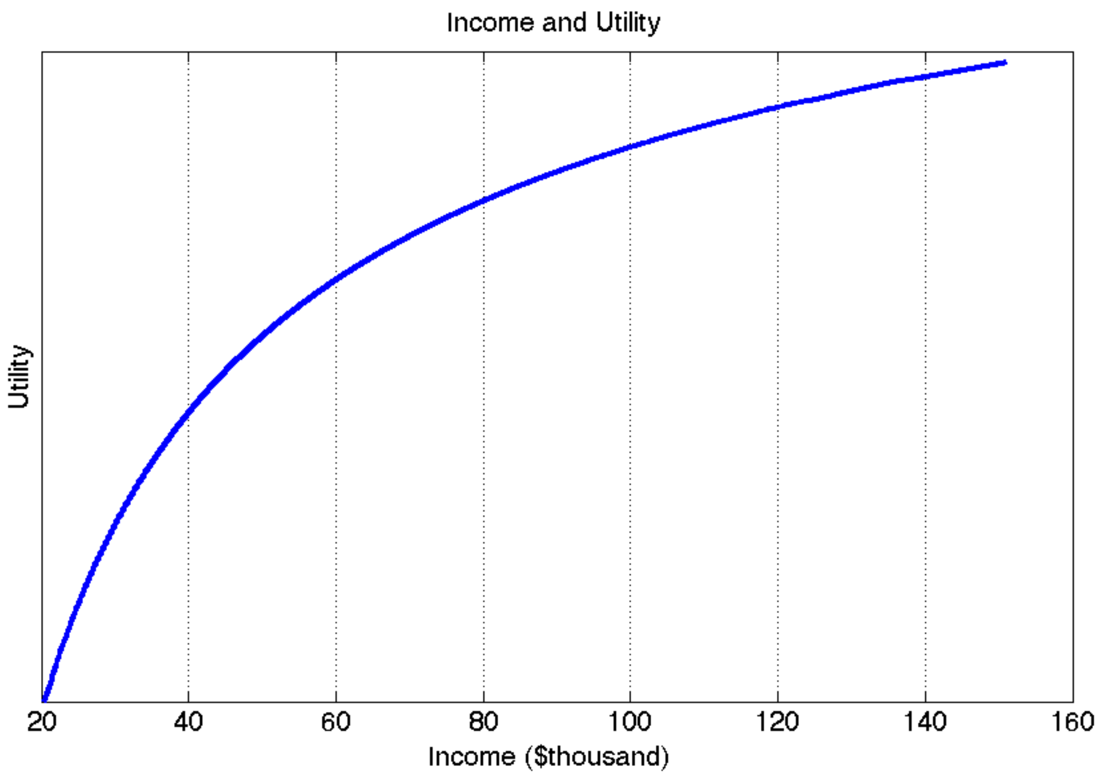

The horizontal axis indicates the total income to be received in a year for consumption over the subsequent twelve months. The assumption is that there is no other source of income for that period. Each scale is arithmetic (not logarithmic).

Note that there are no values on the utility axis. As we will see, this is because in a particular sense, they don't matter. More on that later.

As shown, the slope of the utility curve decreases as income increases. We define the slope as the *marginal utility* of income. Formally, it is the first derivative of the utility function at each point. Informally, it is the rate of change of utility per unit change in income when the latter is very small.

The next figure shows the marginal utility in this case. Note that the greater the income, the smaller the marginal utility. In this case it decreases at a decreasing rate as income increases. Again specific values for marginal utility are not shown because, in a particular sense they do not matter.

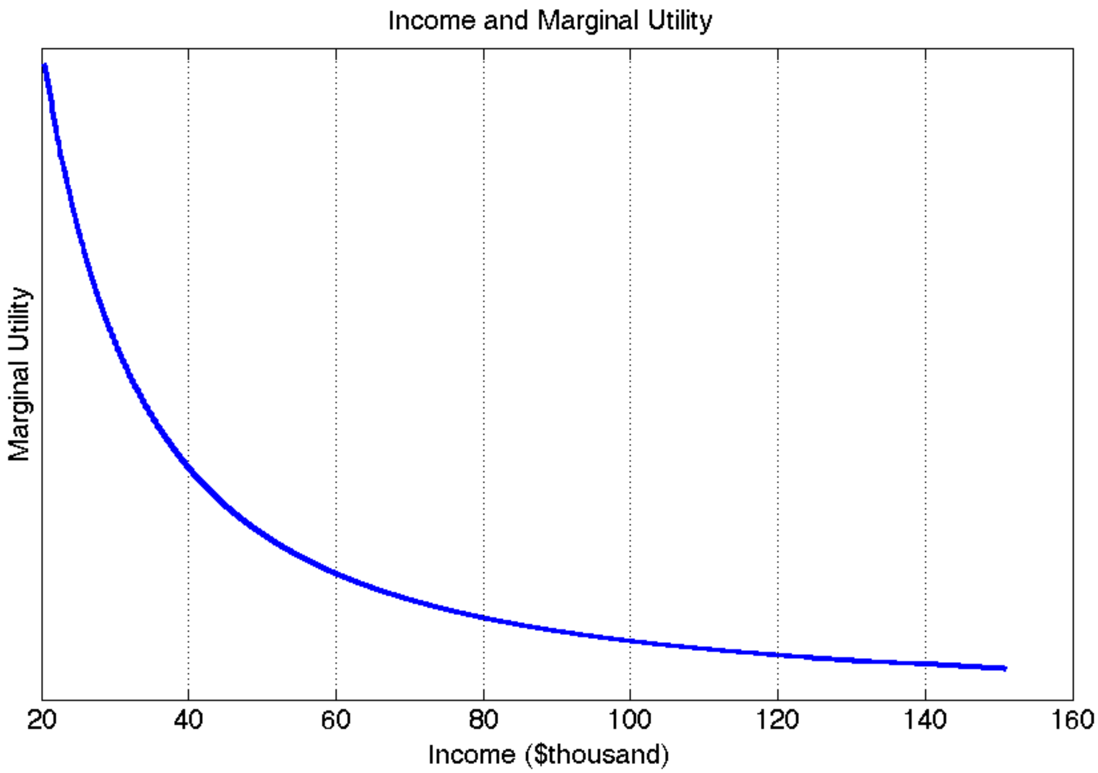

The key aspect here is that marginal utility decreases as income increases – the greater the income, the smaller the additional utility that would be derived from an additional dollar.

Expected Utility

When future income is uncertain, so too will be future utility. But it is possible to compute *expected utility*, obtained by weighting the utility of each possible level of income by its probability, then summing the products. And there is an argument for the assertion that when choosing among alternative distributions of future income, it makes sense to select the one with the greatest expected utility.

Here is the argument in a setting with one future period. Consider a *standard gamble*, with two possible future incomes: one very low, the other very high. In this case, let's say $20,000 or $100,000. Now, ask the recipients to consider a future income of $50,000, then think about a gamble in which there is a *u* percent chance of getting $100,000 and a *1-u* chance of getting $20,000. For what value of *u* would they consider the gamble as desirable as $50,000 guaranteed? Call the answer *u(50)* the *utility* of $50,000 and plot it as a point on the utility function (with $50,000 on the x-axis and *u* on the y-axis). Repeat the question for another possible income, say $60,000, then plot the answer as *u(60)*. Continue until there is a utility curve.

Yes this is whimsical and the sort of approach that behavioral economists love to use as example to ridicule traditional economists as failed mathematicians. But let's stay with it for a bit longer.

Now consider an income strategy that will provide a 50% chance of an income of $50,000 and a 50% chance of an income of $60,000. The first outcome is considered as good as a chance of *u(50)* of winning the top prize (as opposed to the bottom prize). The second is considered as good as a chance of *u(60)* of winning the top prize. The strategy is thus as good as a 50% chance of a *u(50)* chance of winning the top prize and a 50% chance of a *u(60)* chance of winning the top prize. But this is equal to a chance of [0.5*$u(50)$+0.5*$u(60)$] of winning the top prize. And this expression is what we have called the expected utility of the uncertain outcome. Thus, the greater the expected utility of an uncertain prospect, the better it is. And the argument holds for cases with more outcomes and different probabilities of those outcomes. Q.E.D. (*quod erat demonstratum*).

In principle, a financial advisor could go through this kind of exercise with a client or pair thereof, trying to laboriously tease out an appropriate utility function. Of course this seems unlikely, to say the least. But advisors often ask clients to answer a series of questions on a “risk questionnaire” to assess their willingness to take risk in pursuit of higher expected future incomes. Examples can be found on internet-based advisory services. To be kind, one must say that many of these questionnaires are based on little or highly questionable research. But each implicitly attempts to measure the clients' utility function, then recommends an investment strategy that will maximize expected utility using that function. Advisors and clients rarely think in such terms, of course, or do so very informally. Here we will show how to do so formally, then reverse the process, showing how one can analyze a retirement strategy to find the preferences for which it is most suited (the procedure that we will primarily use henceforth).

First-order Conditions for Maximum Expected Utility

Throughout the remainder of this book we shall use the letter *y* to symbolize income (this follows a tradition in some of the economics literature, for reasons mostly unknown). We will also use the term "income" to mean *real income* while avoiding the incessant repetition of the adjective. With these conventions in place, consider a strategy with a (real) income of y_i to be received a year hence in *state i.*

For generality we will consider each cell in our matrices a *state.* Thus each state represents a specific year in a specific scenario. We start with a one-period case in which each state represents income at a particular time in a specific scenario. As indicated earlier, we focus on income to be received a year hence for consumption in the following 12 months.

Let the marginal utility of that income be $m(y_i)$. Let the probability of state *i* be π_i. The contribution of y_i to expected utility will then be $\pi_i m(y_i)$. Finally, let p_i be the price today (present value) of $1 of income in state *i*. The marginal expected utility per dollar of cost for y_i will thus be:

$$\frac{\pi_i m(y_i)}{p_i}$$

Now, imagine a strategy for which the marginal expected utility per dollar is smaller in one state (say, state *i*) than another (say, state *j*). This cannot be optimal, since moving a dollar from state *i* to state *j* will decrease the expected utility provided by state *i* by less than it will increase the expected utility provided by state *j*. For the allocation of income across states to be optimal, the marginal expected utility per dollar of cost must be the same for all states. Formally:

$$\frac{\pi_i m(y_i)}{p_i} = \lambda$$

where λ is some constant.

A little re-arranging gives the following simpler version:

$$m(y_i) = \lambda \frac{p_i}{\pi_i} = \lambda PPC_i$$

As long as marginal utility decreases as income increases, the condition for maximum expected utility is thus that the marginal utility of income in each state should equal a constant times $\frac{p_i}{\pi_i}$ which we will call the *price per chance* (PPC) for that state. The set of such equalities (one for each state *i*) is known as the *first-order conditions* for maximizing expected utility..

Of course, one cannot provide a set of income (y_i) values without limits. Each one costs money and the total cost cannot exceed some pre-determined budget. Since we want expected utility to be as large as possible,the entire budget should be spent. We can write this *budget constraint* as:

$$\Sigma\, p_i y_i = B$$

We now have everything needed to find a set of incomes across states that will provide the maximum expected utility subject to a budget constraint. A key step is to construct a function that can compute the cost of the optimal set of incomes for any chosen value of λ . We will construct one for a particular utility function shortly. Here we simply represent the cost associated with a given value of lambda as $C(\lambda)$.

Now consider the following algorithm for finding the value of λ associated with the optimal set of incomes, given the available budget:

1. Select two values for λ : one (λ_{low}) very small, the other (λ_{high}) very large.
2. Find the costs of the associated optimal sets of income, $C_{low}(\lambda_{low})$ and $C_{high}(\lambda_{high})$. If the budget is between them, proceed. Otherwise adjust one or both of the values of λ as needed.
3. Compute a value λ_{mid} midway between λ_{low} and λ_{high} .
4. Find the cost of the associated optimal set of income, $C_{mid}(\lambda_{mid})$.
5. If this is sufficiently close to the budget B, stop. The associated incomes are optimal
6. Otherwise:
   1. If $C_{mid} > B$ set $\lambda_{low} = \lambda_{mid}$
   2. If $C_{mid} < B$ set $\lambda_{high} = \lambda_{mid}$
7. Return to step 2. Repeat until condition 5 is met.

The termination condition in step 5 can be set to any desired degree of tolerance for a divergence between the cost of the strategy and the budget.

Here is the wikipedia entry for this general approach:

> *The **bisection method** in mathematics is a root-finding method that repeatedly bisects an interval and then selects a subinterval in which a root must lie for further processing. It is a very simple and robust method, but it is also relatively slow. Because of this, it is often used to obtain a rough approximation to a solution which is then used as a starting point for more rapidly converging methods. The method is also called the **interval halving** method, the **binary search method**, or the **dichotomy method**.*

The key words here are *simple* and *robust*. And for our purposes the method is fast enough.

Constant Relative Risk Aversion

Now, from the general to the specific. We start with a widely used type of utility function. The defining characteristic is the fact that it exhibits *constant relative risk aversion;* for short: *CRRA*.

The marginal utility for such a function has the form:

$$m(y_i) = a_i y_i^{-b_i}$$

For generality, we have subscripted each parameter with *i* to allow for the possibility of different values for different states. In this formulation, the two parameters a_i and b_i must both be greater than zero for every possible state.

Now consider the first order conditions for maximizing expected utility. For every state i:

$$m(y_i) = \lambda PPC_i$$

In this case:

$$a_i y_i^{-b_i} = \lambda PPC_i$$

Re-arranging:

$$y_i = (\lambda \frac{PPC_i}{a_i})^{-\frac{1}{b_i}}$$

Using this formula, for any given value of lambda it is straightforward to compute the optimal value of y_i for each value of PPC_i. The resulting cost will be:

$$C(\lambda) = \sum p_i y_i$$

Inserting these computations in our bisection algorithm provides the set of incomes that will maximize expected utility in this setting.

Now from the general to a specific case. Assume that one is choosing incomes to be received at the beginning of year 2 (one year hence), to finance consumption over the subsequent 12 months. Assume that the values of a_i and b_i are the same for every scenario. We start with an example with a value of 1.0 for every a_i and a value of 3.0 for every b_i. Thus:

$$m(y_i) = 1\, y_i^{-3}$$

In this case, the function has a constant relative risk aversion of 3.0; more succinctly: CRRA = 3.

The following figure shows the PPCs and optimal income values for a budget of $50,000 using the state prices obtained using the procedure described in the previous chapter with our standard parameters for the risk-free rate of return and the expected return and standard deviation of return for the market portfolio.

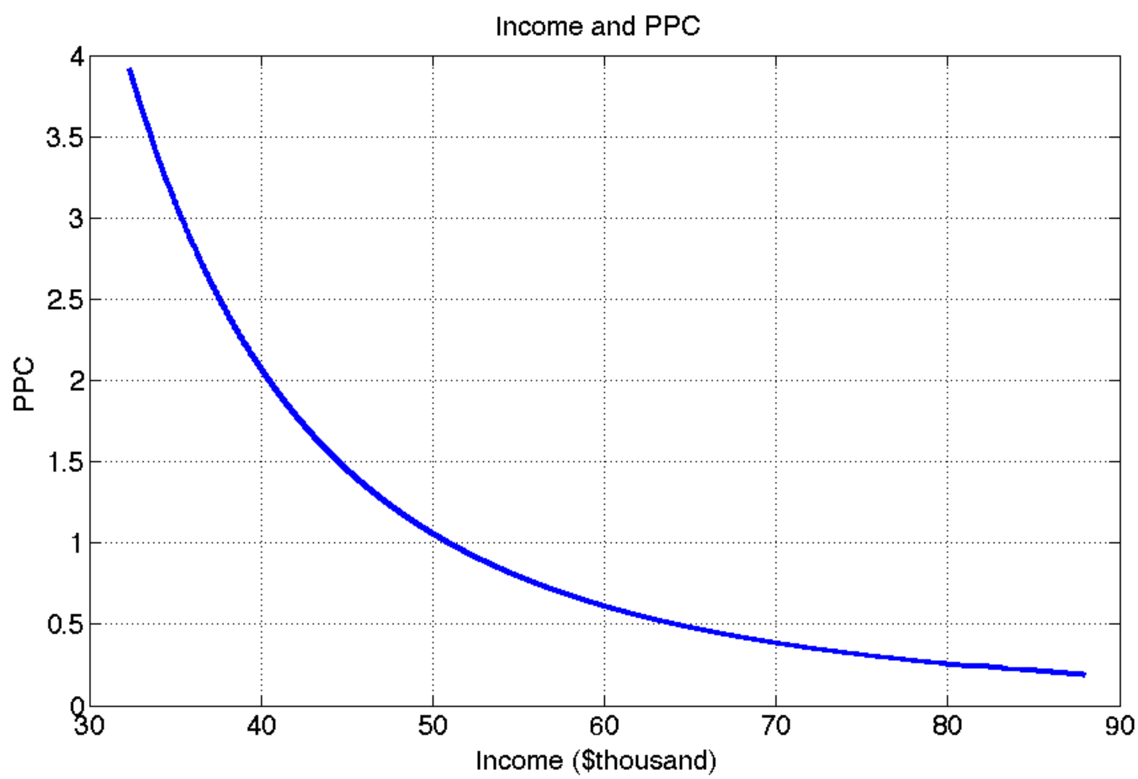

Not surprisingly, the income is a monotonic decreasing function of PPC – the recipients choose to purchase the rights to more income in scenarios when income is cheaper.

As described in the previous chapter, we assume that there is a monotonic decreasing relationship between PPC and the return on the market portfolio. This implies that the relationship between optimal income and market return will be increasing. The next figure, based on 100,000 scenarios, shows that this is indeed the case:

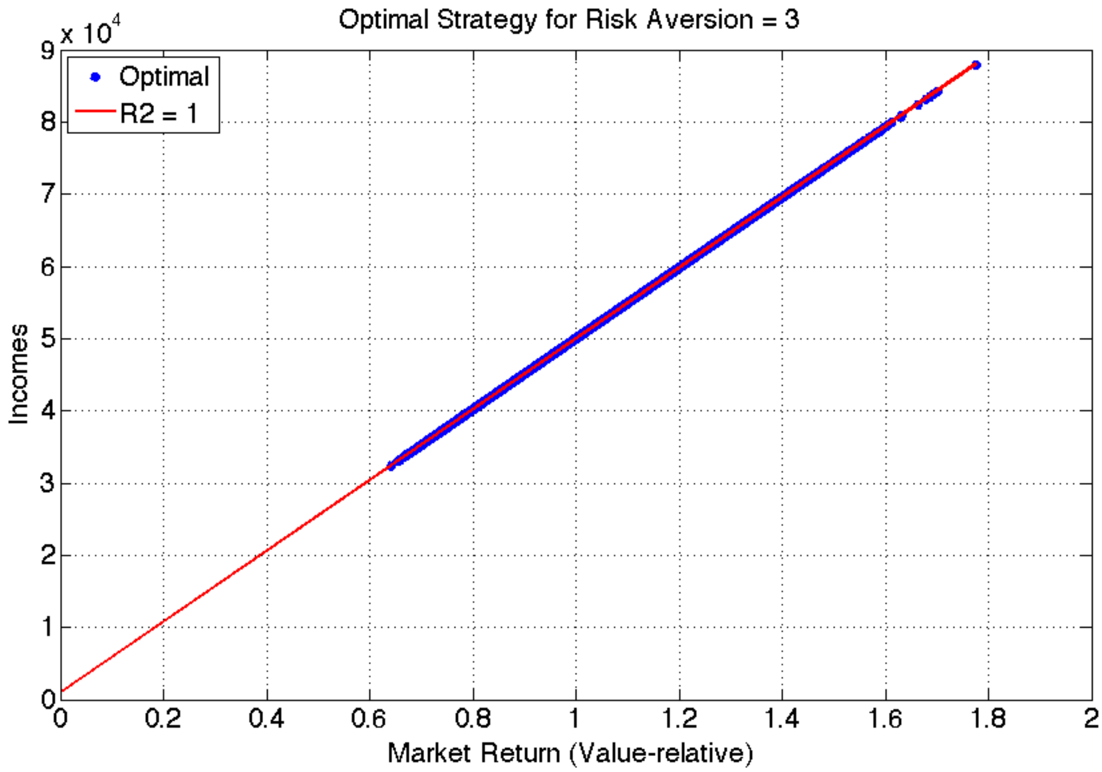

The blue points represent 00,000 levels of income, ranging from over \$30,000 to almost \$90,000 ($9x10^4$). The red line was fit to these points using linear regression. As can be seen, the points fall almost precisely on the line. In fact, the R-squared value representing the degree of fit equals 1.00 (rounded to two decimal places). Moreover, the line passes very close to the origin. The economic implication is clear. The optimal investment policy is to place almost the entire \$50,000 in the market portfolio, hold it for one year, then use the proceeds to provide income in year 2.

Consider now a utility function with a CRRA of 4.0, indicating that the investor or investors have a greater degree of risk-aversion. The next figure shows the relationship between optimal income and the return on the market portfolio.

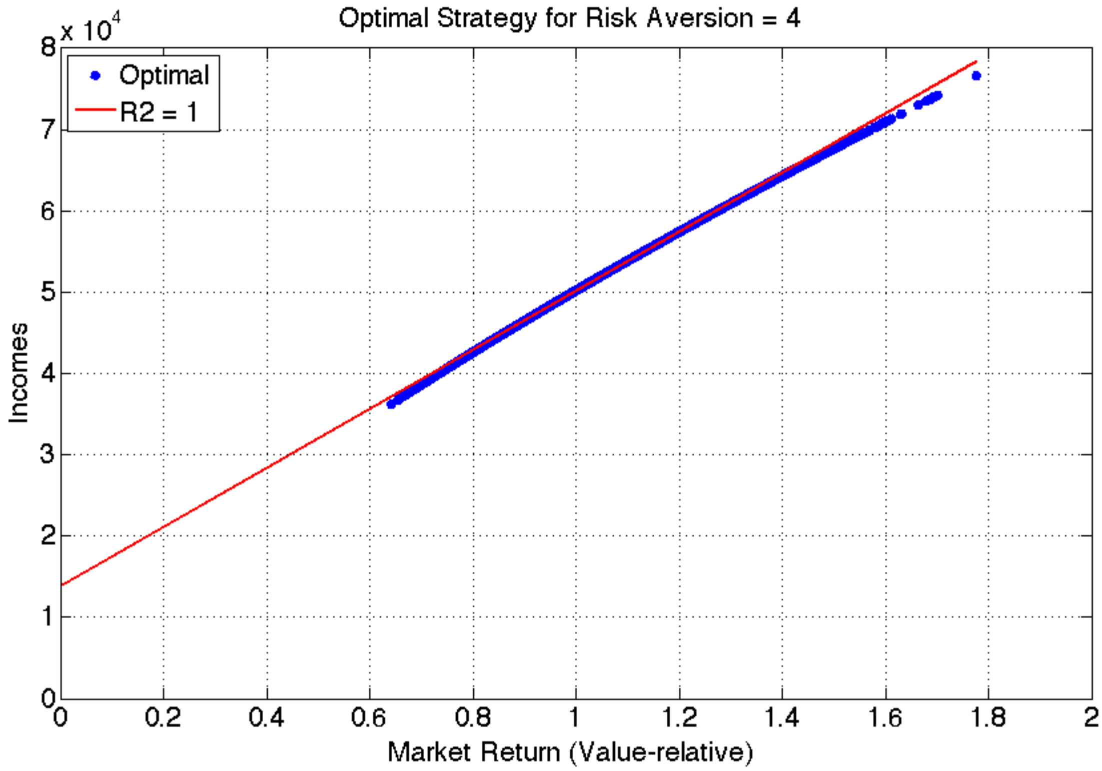

Note that the income levels plot on a curve that increases at a decreasing rate, although the changes in slope are slight. The line through the points fits very well, with an r-squared value of 1.00 (rounded to two decimal places). This may seem strange, but is due to the fact that the great majority of the scenario points lie in the middle of the range and every point is given equal weight when fitting the line.

The most important element in this case is the fact that the intercept of the line on the y-axis is approximately $14,000. This is the amount that would be received if the market portfolio ended up worth nothing at all. This could only happen if $14,000/1.01 had been invested in the risk-free asset, since we are assuming that such an asset provides a real return of 1% per year. The overall strategy is thus very similar to an investment of slightly less than $14,000 in TIPS and the rest (a bit more than $36,000) in the market portfolio.

The next figure shows the results for even more risk-averse investors (with a CRRA of 10.0).

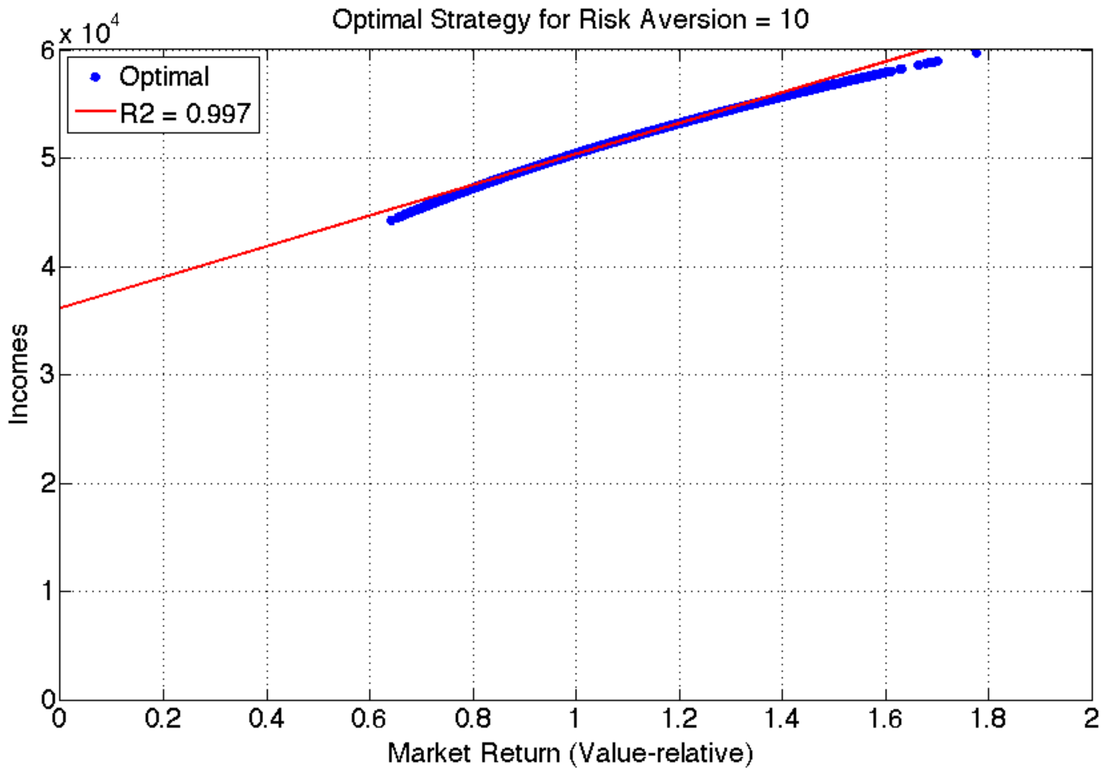

Note that the relationship is still very close to linear (with an R-squared value of 0.997) and the investment in the risk-free asset is even larger (roughly $37,000/1.01).

Finally, consider more adventurous recipients with a CRRA of 2.0. The figure below shows their optimal strategy.

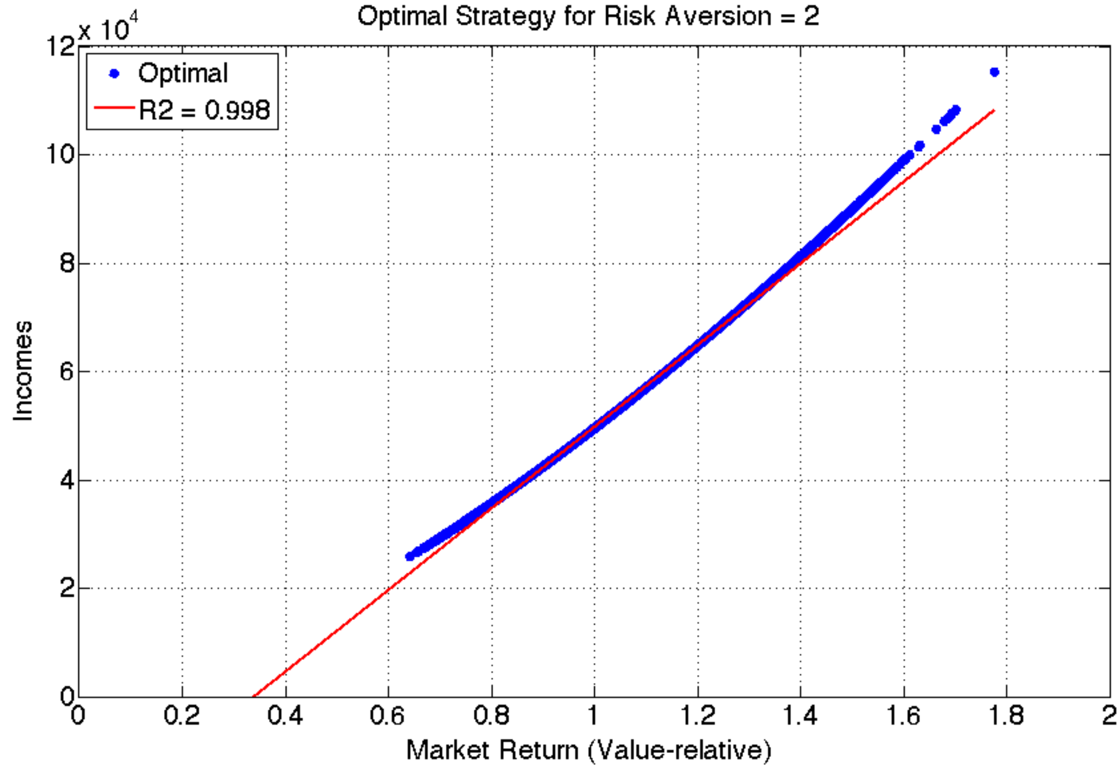

Here there is curvature of the opposite type. Again the regression line fits well, although a small minority of extreme outcomes lie above it. But the line intercepts the x-axis instead of the y-axis. What does this mean? If the graph had been extended to negative values on the y-axis it would have shown that if the market ended up worth nothing, the income would have been negative (roughly -$25,600). This, in turn implies that the strategy involves initially borrowing $25,600/1.01 at the risk-free real interest rate, then investing the original $50,000 plus the proceeds of the loan in the market portfolio. In financial parlance, the investors would use *leverage* to purchase more risky securities than they can afford without a loan. If the market does well, this will provide a large income indeed. But if there is a crash, the net income after repayment of the loan could be small or even negative. In our case, there are no scenarios in which income is in fact zero or negative. But with extreme leverage, this could happen (although it is unlikely that a lender would provide funds in such a situation).

Market-based Strategies

In each of our prior examples, the optimal investment strategy was one in which income was exactly or almost a function of the total return on the market portfolio. In the previous chapter we called any such approach a *market-based strategy,* with the definition expanded slightly to require that income be a non-decreasing function of the total return on the market portfolio.

In our setting, to obtain maximum expected utility with any CRRA utility function requires the adoption of a market-based strategy. This may seem due to the fact that we have only two assets – the market portfolio and a risk-free asset. But it is a far more general result. The key ingredient is the necessity that in order to maximize expected utility, income must be a decreasing (or non-increasing) function of price per chance. And in our view of capital market equilibrium, price per chance is a decreasing function of the return on the market portfolio (that is, a portfolio in which all risky assets held in market proportions). It follows that anyone with a CRRA utility function wishing to maximize expected utility should adopt a market-based strategy.

The result is even more general. All that is required is that the investors' marginal utility function be a decreasing (or non-increasing) function of income. A market-based strategy will still be optimal, although the relationship between income and market return may have some flat spots, as we will see.

The Representative Investor

It may seem strange that in our example, the optimal strategy for a CRRA investor with a risk-aversion of 3.0 is to invest almost all assets in the market portfolio. But there is a simple reason. Recall the equation for the pricing kernel from the previous chapter:

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

This has the same form as a CRRA marginal utility function. And with our standard assumptions about the real risk and return of the market portfolio and the real return on the risk-free asset, the coefficient *b* was 2.9428 – very close to 3.0. Note also that the marginal utility for an investor with a CRRA utility will plot as a straight line in a log-log graph which will exhibit constant elasticity, as does our pricing kernel.

Now, imagine a single-period market with a single investor having a CRRA utility function with a risk-aversion of 2.9428. State prices for money received at the beginning of year 2 would be equal to those in our matrix. In effect, market prices are set as if there were a *representative investor* with a CRRA utility with a risk-aversion of 2.9428. Of course investors have diverse preferences, so the pricing kernel represents (speaking loosely) a kind of average of their marginal utilities. But in our world, the net result for asset pricing is the same as if there were just one "representative" CRRA investor with this risk aversion.

This shows (1) why a CRRA investor with a risk-aversion of 3.0 should invest almost all funds in the market portfolio, (2) why a CRRA investor with a greater risk-aversion should hold both the market portfolio and a risk-free asset, and (3) why one with lower risk-aversion should lever up the market portfolio. When considering an investment policy, a key question is whether you are more or less risk-averse than the representative investor.

It is also important to remember that security markets are not democracies, in which each person has one vote. In the determination of asset prices, rich investors have more votes than poor investors. When considering an investment strategy it is important to compare oneself with an image of a relatively wealthy individual or a financial institution. Look to Wall Street, not Main Street. With this view in mind, it seems reasonable to conclude that most retirees should invest more conservatively than the "representative investor".

Utility and Marginal Utility

It may seem strange that we have found a way to maximize utility without ever defining it. And it is time to clear up the mysterious earlier statements that the values on the vertical axis of a utility graph don't matter.

First, assume that the utility of income is represented by a function $U(y)$. Now imagine adding a positive constant to each utility value, so that the new utility function is $k+U(y)$. Clearly, the derivatives (marginal utilities) of this function will be the same as those of the original function. Hence the set of incomes that will maximize the expected value of the new function is the same as the set that maximized the old one. To find the maximum expected utility, we need only the derivatives (marginal utilities).

The second argument may only apply to a single-period case in which the constant term for the utility is the same for all the states. Imagine that the original utility function is multiplied by a different (positive) constant so the new utility function is $kU(y)$. Now each marginal utility will be a constant times its previous value. But when the optimal set of incomes is computed, this will give the same result; only the value of λ will differ. In our formulation, if all the constant (*a*) terms are the same, they could all be replaced with 1.0. But, as we will see, in a multi-period setting it may be important to have constant (*a*) values that differ by year and personal state, and the magnitudes of such differences will affect the amounts of optimal incomes in different states.

The earliest economists devoted a great deal of thought to the concept of utility. In 1789 Jeremy Bentham wrote:

> *By utility is meant that property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness, (all this in the present case comes to the same thing) or (what comes again to the same thing) to prevent the happening of mischief, pain, evil, or unhappiness to the party whose interest is considered.*

We need not attempt to measure anything this profound. Instead we focus on marginal utility as a relatively pragmatic measure of willingness to trade consumption in one period or state for that in another. We leave the task of actually measuring happiness to moral philosophers, neural scientists and others.

Revealed Preference

We know from the first order conditions for optimization that if a set of incomes is optimal for a particular marginal utility function and budget, the marginal utility of the income in each state will equal a constant times the PPC for that state, that is:

$$m(y_i) = \lambda PPC_i$$

This allows us to infer an investor's utility function from the properties of a chosen strategy (on the assumption, of course, that the choice was made appropriately). In a diagram with *PPC* on the vertical axis and income (*y)* on the horizontal axis, we simply remove the numbers from the vertical axis and interpret the diagram as the revealed preference of the investor. If the relationship is not monotonic, we can replace it with a monotonic function with the same cumulative probability distribution (the cost-efficient alternative discussed in the previous chapter), then infer an associated marginal utility function, although this assumes that the investor intended to use the efficient version (which requires a substantial leap of faith).

More generally, we can evaluate any market-based strategy by attempting to determine if the *revealed preference* (marginal utility) is consistent with the investor's actual preferences. Or we can say that a particular market-based strategy is optimal for an investor with the associated marginal utility. An example based on behavioral research illustrates the point.

Reference Levels of Income

An investor's marginal utility curve could take one of many possible forms. Some alternatives, like the CRRA version, are continuous. But not all. In numerous experiments, cognitive psychologists have found that when presented with simple decisions involving uncertain monetary rewards, individuals tend to make choices that are inconsistent with continuous marginal utility curves. In particular, people seem to focus on a *reference point* (such as their current wealth or income) and to feel that a small loss from that point would be considerably worse than an equally small gain would be good. For example, a subject might require a 50% chance of winning $2 to be willing to take a 50% chance of losing $1. Such preferences are a key part of *prospect theory,* developed by Daniel Kahneman and Amos Tversky. In their formulation, the marginal utility of a value slightly greater than the reference point is considerably less than the marginal utility of a value slightly below it. In effect, the utility function has a *kink* at the reference point and the associated marginal utility curve is discontinuous at that point.

The figure below shows the associated marginal utility curve. In principle, it is discontinuous at the reference level of income (here, $50,000) but we have drawn a vertical line. For some purposes it is useful to think of this segment as having a very slight slope so that the relationship between marginal utility and income will be monotonic.

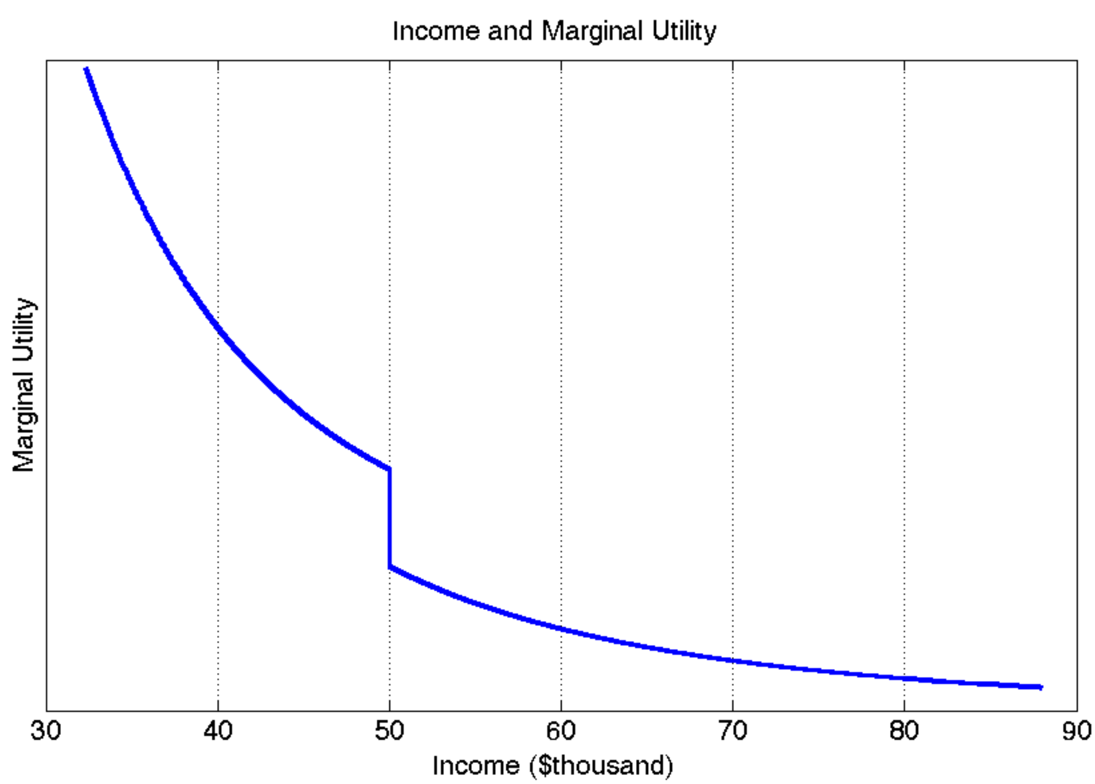

It is easy to guess the optimal relationship between the return on the market portfolio and the optimal income for an investor with this sort of marginal utility. Within some range of returns on the market portfolio, the income will be the same (or virtually the same) and equal to the reference income. The following figure provides an example in which the marginal utility function has constant relative risk aversion above and below the reference point of $50,000.

In later chapters we will see examples of financial strategies that produce flat or nearly-flat sections in this sort of diagram, presumably designed to appeal to retirees for whom reference levels of income are important.

Tranches

As our earlier examples showed, it can be relatively easy for an investor with constant relative risk aversion to obtain a market-based strategy that comes close to maximizing expected utility. One would just allocate some money to the market portfolio and the rest to a risk-free security (including the possibility that the allocation to the latter is a negative number). But what is the investor with the kinked utility curve to do? As the previous figure showed, the optimal relationship between income and the market return goes up, then is flat, then goes up again.

Those who design complex financial instruments term an investment with payoffs of this type a “Travolta”, after the one of the dances performed by the actor John Travolta in the film “Saturday Night Fever”. To wit:

Interestingly, there are also images online in which the positions of his arms are reversed, but there seems to be little interest in a financial product that goes down as the market rises with a flat space between.

While clever financial engineers can and do create products with such payoffs using instruments called options and other derivatives, such exotica can involve large obvious or hidden costs and/or the possibility of partial or complete default. But if demand were large enough, there is a simpler and safer way.

Consider an investment fund that invests $60,000 in the market portfolio and $20,000/1.01 in the risk-free asset. The value of the portfolio a year hence will depend on the return on the market portfolio, as shown by the red line in the figure below.

Now assume that the fund issues two classes of shares. Holders of Class A are promised payments a year hence that will depend on the market return in the manner shown by the blue curve. Class B holders are promised payments equal to the remainder of the fund after the Class A shareholders have been paid. Graphically, the amount paid to Class B will equal the vertical difference between the red curve and the blue curve. In industry parlance, the total portfolio return has been divided into two *tranches*.

Of course there will be expenses, including compensation for the financial firm that created the fund. But there should be no default risk. The TIPS and market portfolio securities can be put in a safe, inspected periodically by some enforcing agency, then cashed out at the termination date (here, a year hence). And this sort of fund can be created for any desired payment schedule (the blue curve) by holding some combination of the risk-free asset and the market portfolio. The total portfolio value (red curve) will be a linear function of the market return and should have enough invested in the market portfolio so that the residual (Class B) share return is also a non-decreasing function of the market return, as in this example.

The argument can be extended to cover multiple years and more complex functions of the market portfolio (market-based strategies) for Class A. Given sufficient demand, almost any type of desired strategy can be accommodated.

Note that investors who want a Travolta payoff (Class A) must find other investors who want the complement (Class B). More generally, markets must clear. If one investor wants to be protected from market declines, some other investor must be overly exposed to them. Security prices will adjust until there is supply for every demand. In principle, the prices of our Class A and B shares will be determined by our pricing kernel and both classes will find willing buyers.

In my 2007 book (*Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice, p. 167-168*), I used the term *m-shares* to describe tranches such as these in which the underlying asset is a market portfolio. By the beginning of 2017, neither the practice nor the terminology had caught on. But hope springs eternal.

As we will see, several retirement income strategies provide payments that are non-linear functions of the return on the overall market. However, their returns are typically not cost-efficient, suggesting that there might be sufficient demand for new financial products with payment tranches based on market returns – a subject to which we will return in later chapters.

Utility Dependent on Time and Personal State

Many utility analyses focus on a single time period in which utility and marginal utility are solely functions of income or consumption. But our focus is on decisions that involve income in multiple time periods, with possibly different personal states in each time period. This substantially increases the difficulty of finding a strategy that will maximize expected utility, given an overall budget.

Formally ,the problem is unchanged and can be solved using the algorithm described earlier. The key step is to select a set of a_i and b_i values for utility functions taking into account both the year and personal state for each of the possible scenarios and personal states. Computationally, we need to fill two matrices, each the size of our previous matrices, with scenarios as rows and years as columns. The first will have an a_i value in each cell and the latter will have a b_i value in each cell. Within a column (year) all the a_i values will be the same for all scenarios with a given personal state, as will all the b_i values. As we will see in later chapters, many people make choices consistent with preferences in which for each personal state the values of one or both parameters vary from year to year. Thus many retirement income scenarios appear to be optimal for investors with marginal utility functions that depend on both the year income is to be received and the personal state of the recipients at the time.

This said, the effect of a personal state on utility is likely to be only partially taken into account when many retirees choose a retirement income strategy. Consider first the difference between (a) states 1,2 and 3 (when one, the other or both retirees are alive) and (b) state 4, the first year in which neither is alive and state 0 for each year thereafter. As the aphorism says about money, “You can't take it with you”. But you can, of course leave it to your estate or spend it all before you die. To put it rather crudely, for many if not most people, the utility for money in “alive” states differs from that in “dead” states. If so, the amount of income should differ as well. But there is no practical way for an individual or pair thereof to accomplish this on their own since under most circumstances mortality is a probability, not a certainty.

But there is a way to arrange for income to be contingent on personal states: create a mortality pool with a number of others, so that *probabilities* can become *frequencies*. Thus I might have a probability of dying next year of 10%, but in a pool with thousands of people of my age and sex, it is almost certain that close to 10% will die. Creating and servicing such pools is the task of private insurance companies and, in most countries, governments. We will examine private annuities and government social security programs in detail in later chapters. At this point it suffices to point out that many retirees with discretionary funds that could be used to supplement social security payments choose not to purchase annuities. Possible reasons are: (1) they equate the expected utility from their own consumption with that of consumption by their heirs, (2) they wish to avoid the cost of insurance company fees or (3) they have not seriously considered the economics of the alternatives. More on this anon.

Time-separable Utility

Thus far we have assumed that the expected utility of a retirement income strategy can be determined from utilities computed for each scenario/year cell in our matrix. In particular, the utility of consumption in year t depends only one's personal state and the year in question. The *sequence* of consumption values does not play an additional role. Formally, let income in years 1, 2 and 3 in a particular scenario be $[y_{1,} y_{2,} y_3]$. If the associated utility can be computed using the three values separately, as in $[U_1(y_1)+U_2(y_2)+U_3(y_3)]$, we say that utility is *time-separable*. But this may not reflect the true preferences of some retirees. Some of the behavioral economics literature suggests that people are may also be concerned with the *sequence* of incomes (recall, for example, the idea of a reference income). In this more general view, utility is a more complex function and should be written as $U(y_{1,} y_{2,} y_3)$.

Here is a simple but informative example. Assume that a retiree has savings placed in three *lockboxes*, each of which is to be used to provide income in a specific year. The riskless rate of interest is 1%, and initially the first lockbox has $100, the second $100/1.01 and the third $100/(1.01^2). Income for the first year (to be provided immediately) will come from lockbox 1. Income for the second year (to be provided twelve months hence) will come from the proceeds of the investment of the funds in lockbox 2. And income for the third year (to be provided in 24 months) will come from the proceeds of the investment of the funds in lockbox 3.

In this setting, the key questions are how to invest the funds in lockboxes 2 and 3. First, it is clear that if both are invested solely in the risk-free asset, income will be $100 in each of the three years (explaining why the initial amounts were chosen as they were). But what if the retiree is willing to take risk in order to increase expected income? We will assume that lockbox 2 is invested in a market portfolio offering our standard lognormal return distribution with an expected real return of 5.25% and a standard deviation of 12.5%.

We focus on three different strategies for investing the funds in lockbox 3, which we will call Strategies A, B and C.

Strategy A will invest the funds for year 3 in the market portfolio for the first year, then in the risk-free asset for the second year. It will then provide precisely the same income in year 3 as in year 2, no matter what the latter may have been. This can be seen in the figure below in which the y-axis shows the probability of exceeding each possible income on the x axis (the manner in which we will choose to summarize income probability distributions, as will be discussed in later chapters).

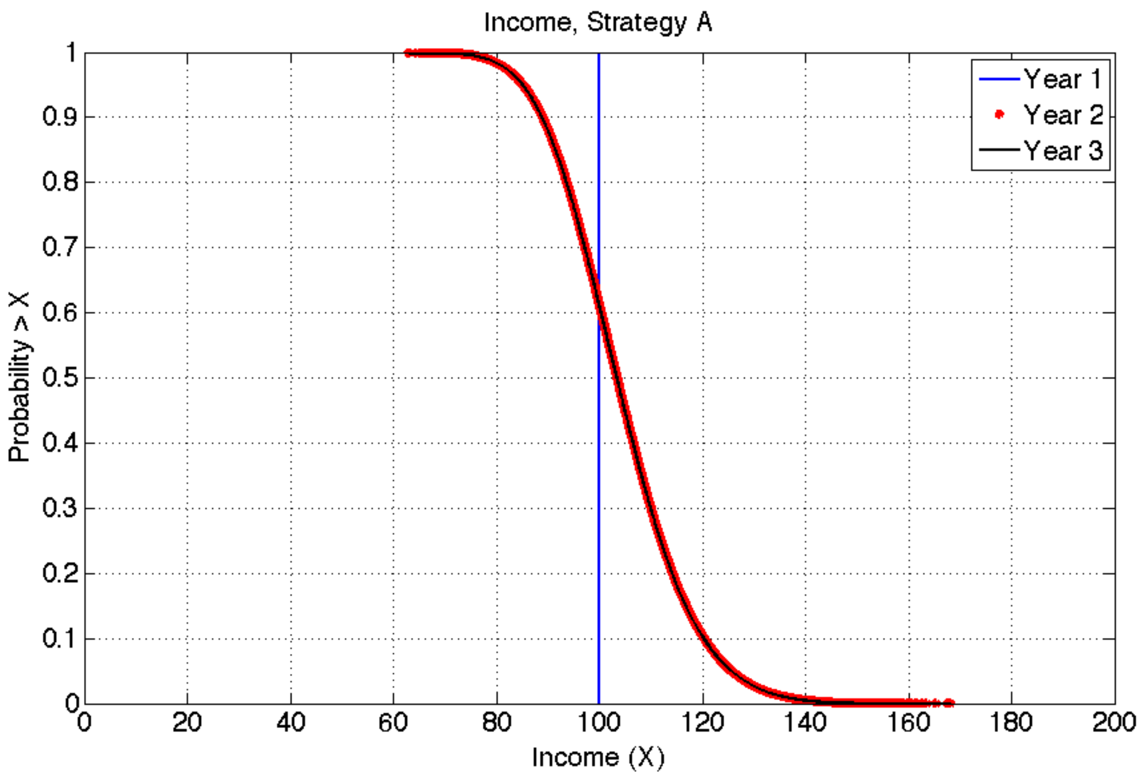

As intended, the probability distribution for income in year 3, when viewed from the present, is the same as that for income in year 2.

Note, however, that with this strategy, when viewed from year 2 the income for year 3 is known with certainty. It will, in fact be exactly the same. This can be seen in the figure below which shows the possible ratios of income in year 3 to that in year 2. In this sense, there is no *sequence risk* with Strategy A.

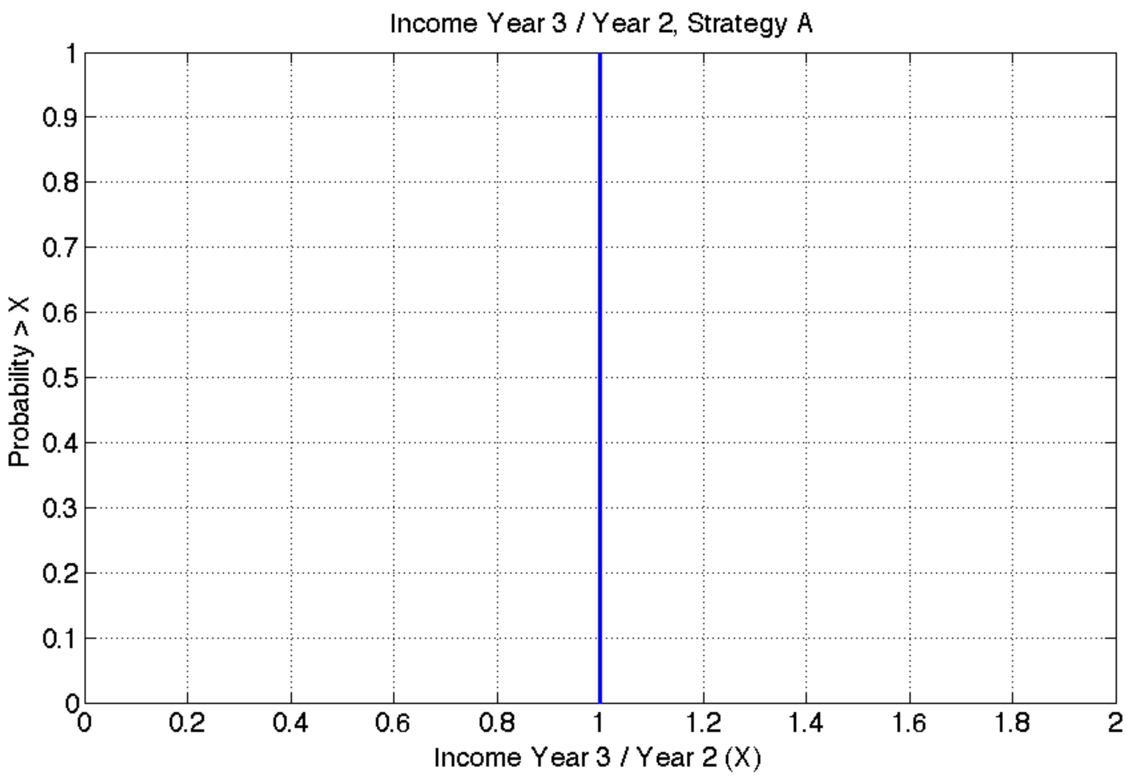

Now let's turn to Strategy B. It differs from Strategy A only in the investment rule for lockbox 3. In this case the initial amount is invested in the risk-free asset for the first year, then the proceeds are used to purchase the market portfolio in the second year. This will give the same distribution of income in year 3 as does Strategy A, which will also be the same as that for income in year 2, as shown below.

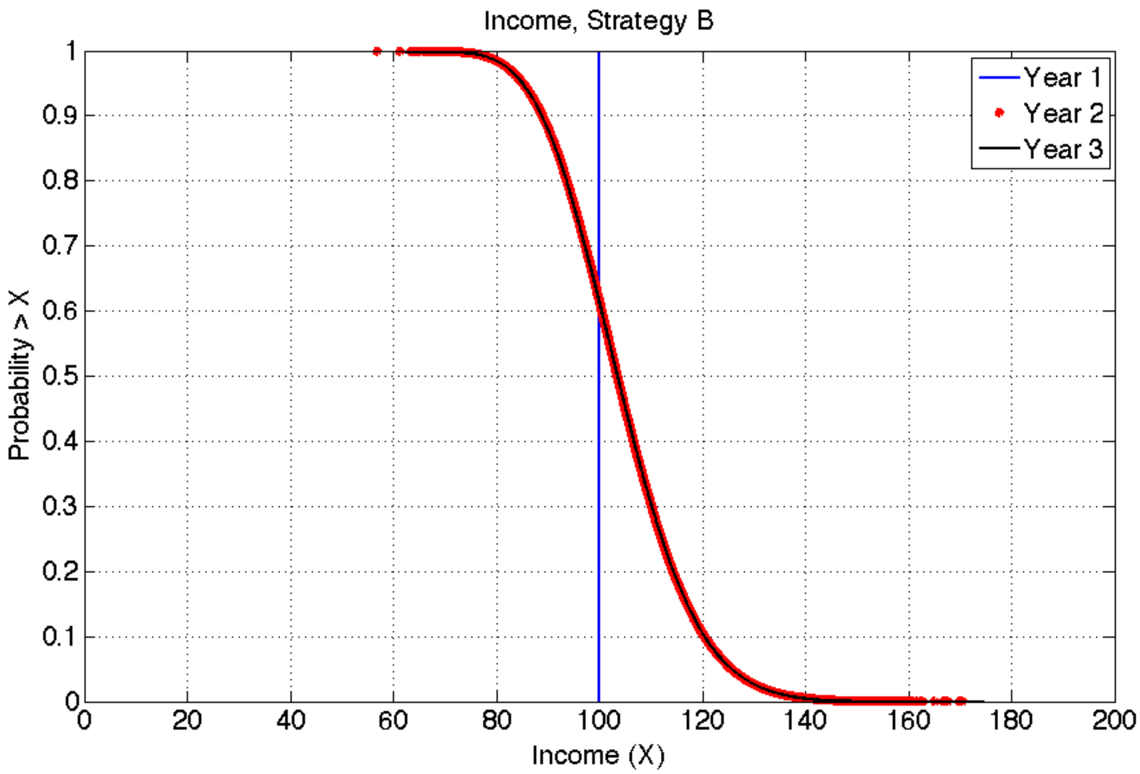

But the year-over-year results will be very different, since the incomes for year 2 and 3 depend on the market returns in two different years. The figure below shows the results.

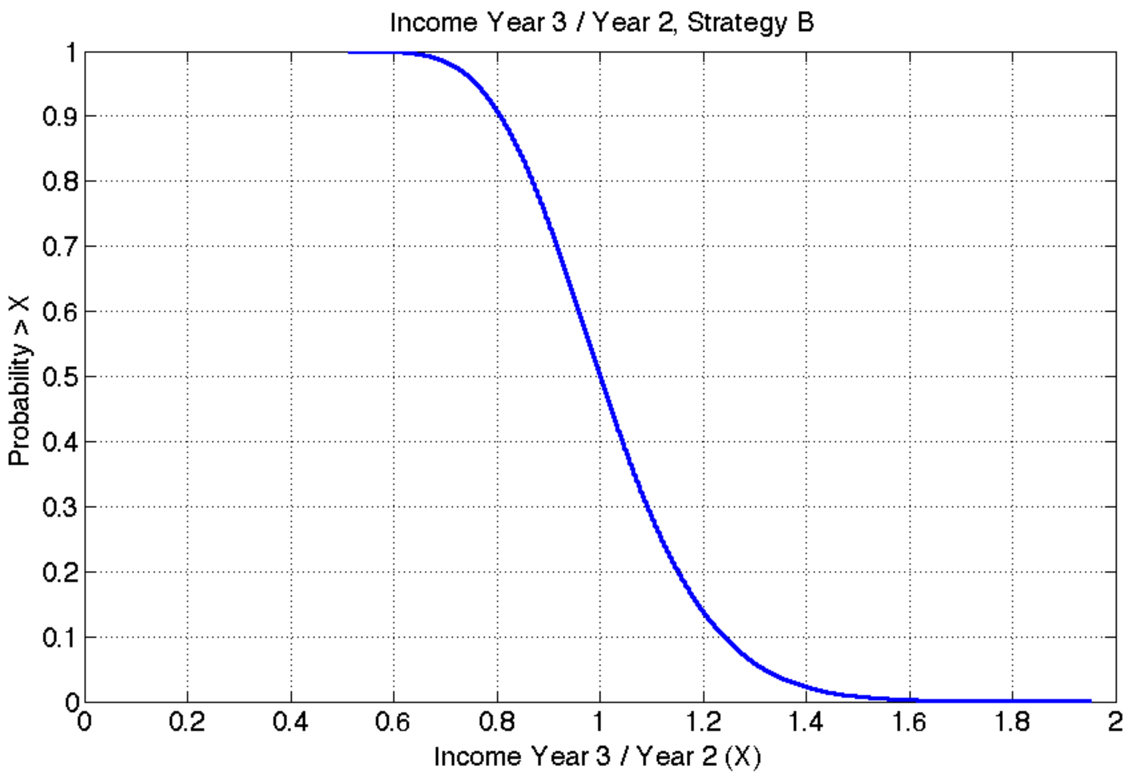

Now, assume that you had to choose between strategies A and B. Each provides the same probability distribution of income for years 2 and 3 as viewed from today. But Strategy A resolves all uncertainty about income in year 3 a year in advance, while Strategy B does not. For an investor with time-separable utility, the two should be equally desirable. But someone concerned with sequence risk could prefer Strategy B.

You might wish to pause at this point to decide which you would prefer. But this example suggests that it is useful for recipients to at least consider the aspect of risk captured in a graph of the range of possible ratios of income from year to year. For this reason we will make it possible to include such information routinely in analyses of alternative strategies.

There is more to be said about this example. Neither Strategy A nor Strategy B is *cost-efficient* in the sense discussed in the previous chapter. Each provides income for year 3 that depends on both the cumulative return on the market portfolio over two years and the particular path taken to achieve that market return. The results are *path-dependent* and there is some alternative market-based strategy can provide the same probability distribution of income at lower cost.

As shown earlier, for a cost-efficient strategy, income is a non-increasing function of PPC and a non-decreasing function of return on the market portfolio. This is clearly not the case for either Strategy A or B. The following figure shows the situation for the Strategy B.

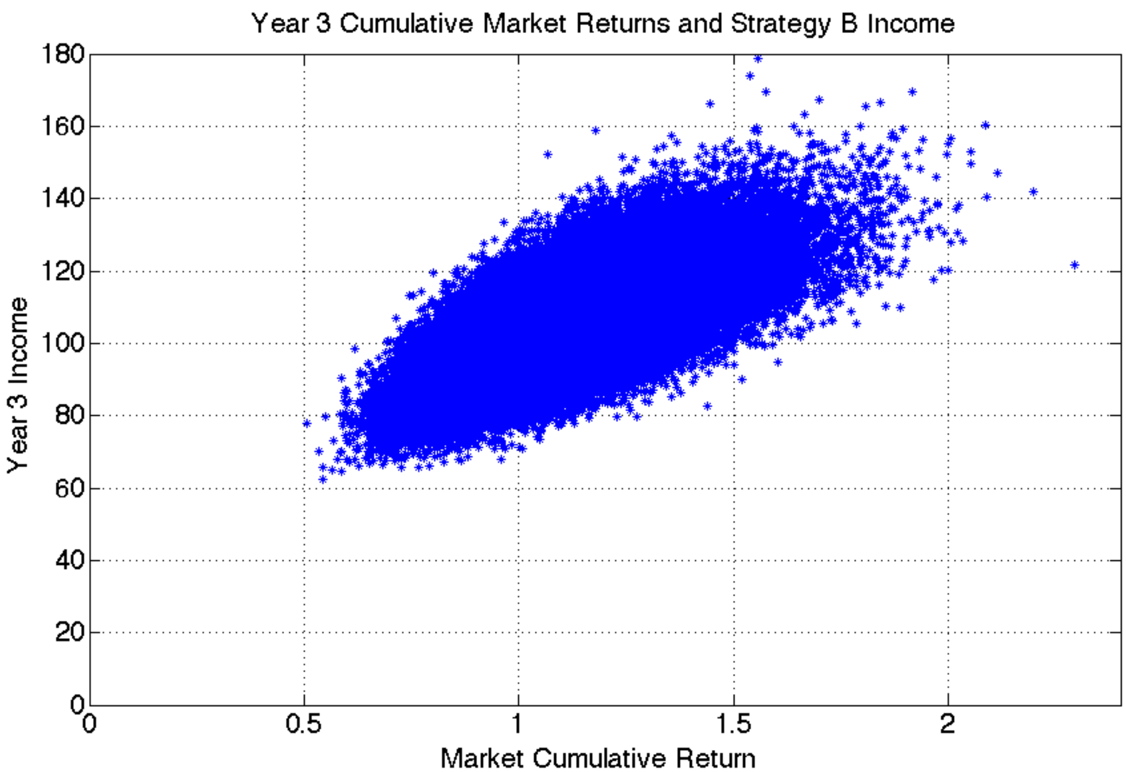

But we know how to create a market-based cost-efficient strategy for year 3 that will have the same probability distribution of income in year 3 as that of Strategy B. The blue points in figure below show the result (obtained by sorting both the market cumulative returns and the year 3 incomes from Strategy B in ascending order).

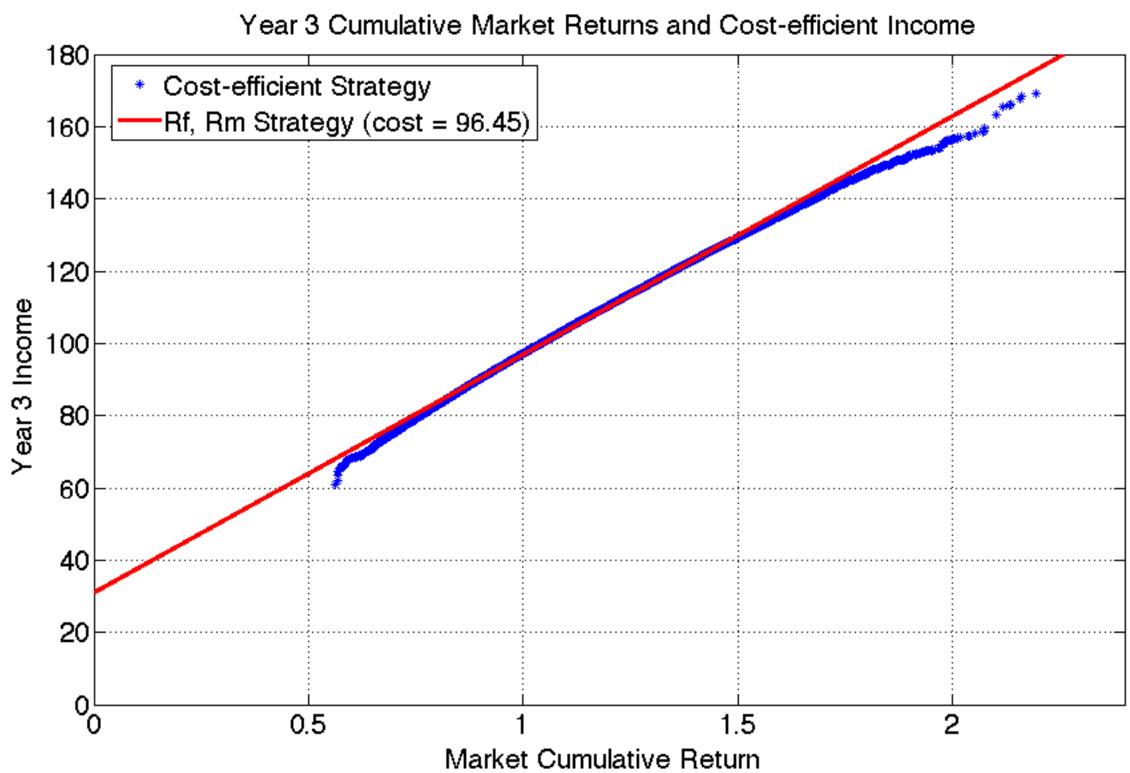

As we before, we fit a straight line to the points using linear regression. The fit was extremely good, with an R-squared value of 0.9985 (since the vast majority of the scenarios were in the middle part of the market return range, where the relationship was almost linear). The points on the line could be achieved with a simple strategy in which lockbox 3 is initially invested in a combination of the risk-free asset and the market portfolio, left alone for two years, after which the securities are cashed in to provide income for year 3.

In this case the regression equation was:

$$y = 31.02 + 66.04\, Rm$$

The economics are straightforward. If the market portfolio were to be worth nothing, the lockbox value would be $31.02. Since this would be provided by two years of compound returns on the risk-free asset the original investment in that asset would be 31.02/(1.01^2) or $30.41. If the final market portfolio were to equal its original value, the portfolio would be worth $66.04 more than the amount provided by the risk-free asset, thus the initial amount invested in the market portfolio would be $66.04. The total cost would thus be $30.41+$66.04, or $96.45. Note that this is lower than the $98.03 cost of lockbox 3 for the other strategies ($100/(1.01^2).

This is our Strategy C. Viewed from today, it provides the same probability distribution of income for each of the three years as do Strategies A and B. This figure makes the point.

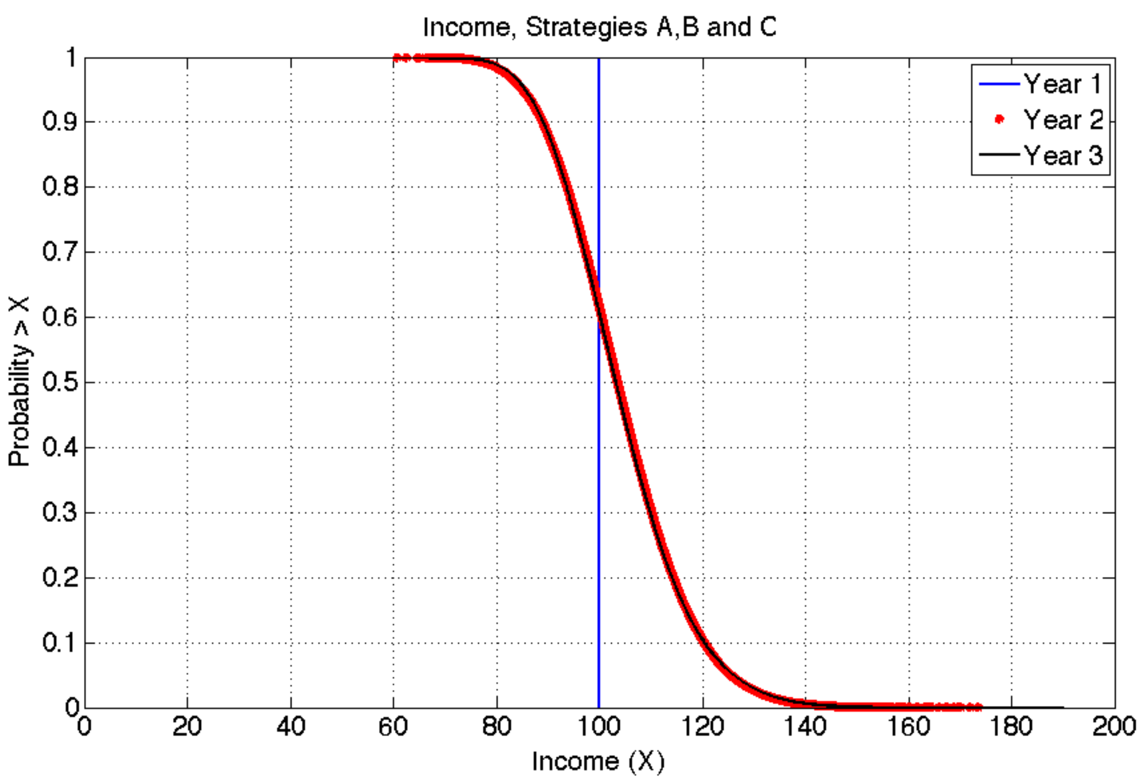

But the strategies provide different probability distributions of the ratio of income in year 3 to that in year 2, as this figure shows.

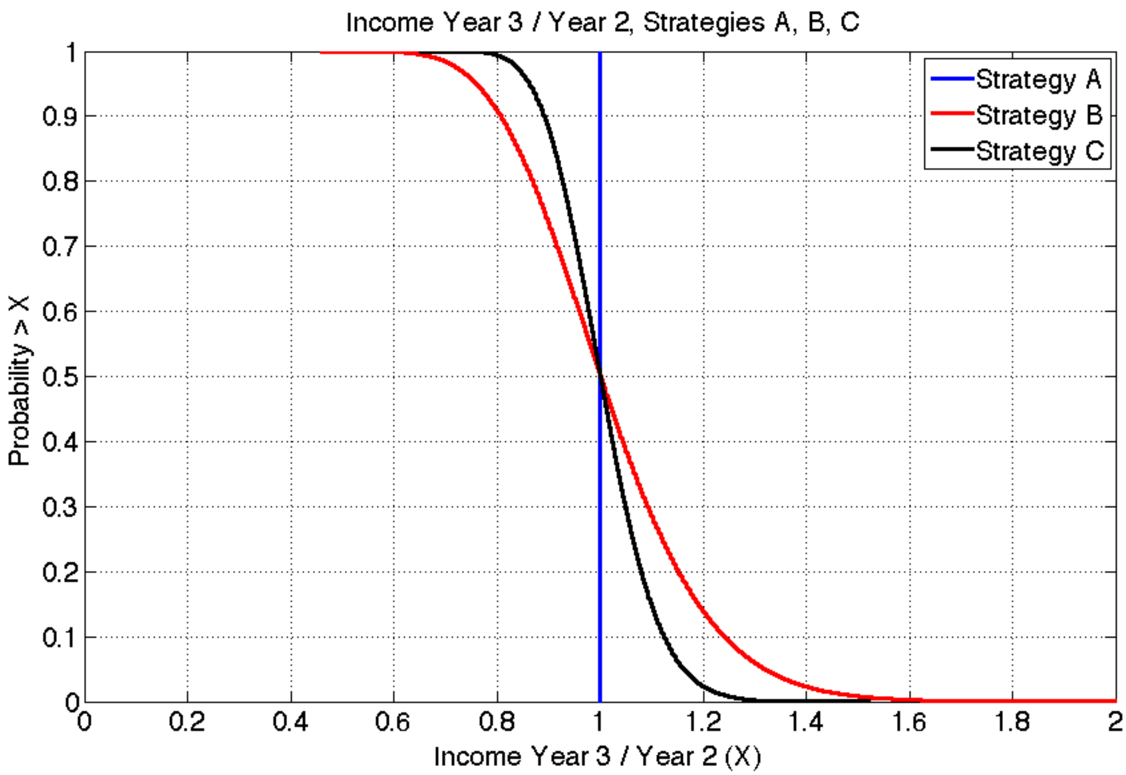

For someone concerned with year-to-year variations in income, Strategy C is better than Strategy B. It provides the same distribution of income in year 3 and is cheaper. Thus it dominates the latter. The choice between strategies A and C may be somewhat harder. Strategy A has less year-to-year variability in the ratio of income in year 3 to that in year 2, but Strategy C is cheaper. Anyone with time-separable utility would choose C. Others could choose any of the three.

This example emphasizes the fact that our concept of cost-efficiency is predicated on the assumption of time-separable utility. For those with such preferences, a cost-efficient strategy dominates one that is not cost-efficient and market-based strategies are to be preferred. Those concerned also year-to-year variability in income might prefer strategies with some path-dependent outcomes. That said, the very notion of declining marginal utility gives preference to strategies without excessive year-to-year income variation. Moreover, the best way to reduce year-to-year income variability may be to choose risk-free investments. We will have much more to say on these issues as we analyze different retirement income strategies.

Implied Marginal Utilities

Before we leave strategy C it is useful to examine the marginal utility functions for which it would be optimal. The figure below shows the relationships between the logarithm of income (on the horizontal axis) and the logarithm of PPC, the price per chance, (on the vertical axis) for each of the two future years.

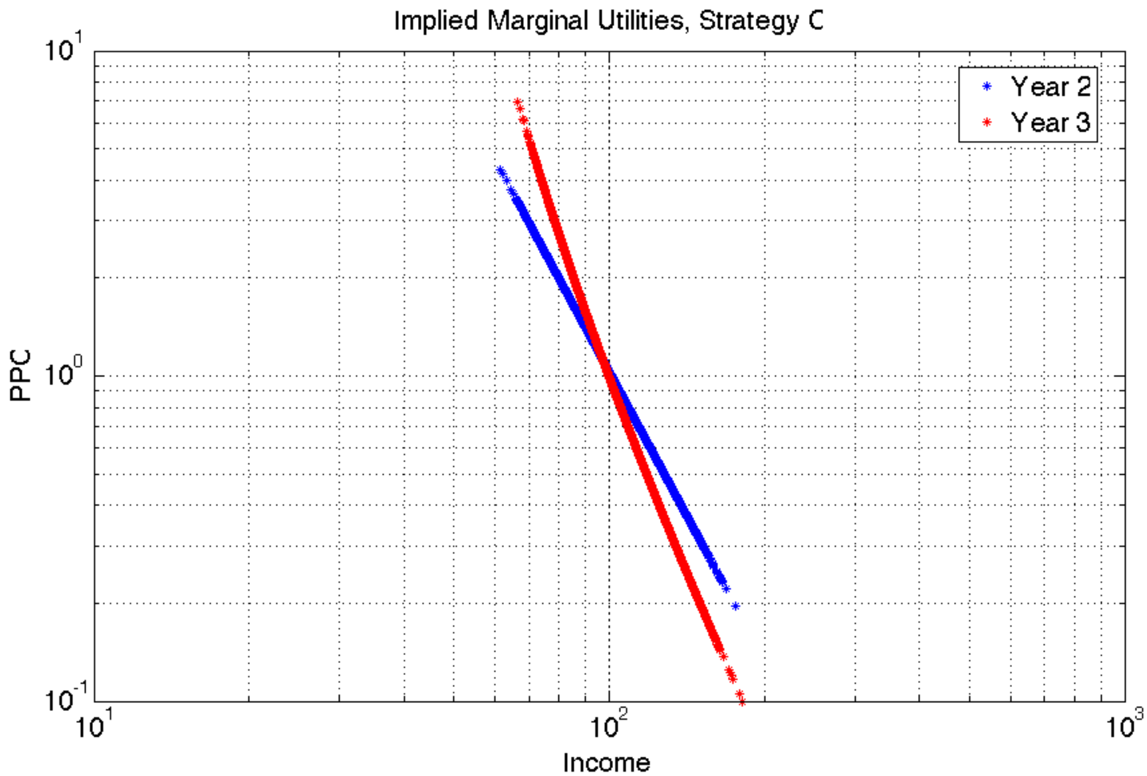

It is not surprising that the function is linear in this diagram for year 2 income, since it is basically a constant times the one-year return on the market portfolio. We know that a straight line in a diagram of this type indicates an implied marginal utility reflecting constant relative risk aversion. The regression line for year 2 has a slope of -2.94, equal to the elasticity of the market return, with an R-squared value to three decimal places of 1.0000. Not surprisingly, the risk-aversion of someone who one invests in the market portfolio for a year is 2.94, equal to that of the market representative investor.

The implied marginal utility function for year 3 is also virtually linear, but steeper, with an implied constant relative risk aversion of 4.21. The points fit the line almost as well, with an R-squared value of 0.9987. This reflects the more conservative decision to invest in both the risk-free asset and the market portfolio. Because there is more variation in the PPC values for periods farther in the future, a decision to keep the range of possible incomes relatively constant from year to year reflects greater risk-aversion, the farther in the future the period in which income is to be received.

It is possible to fit a regression line to any set of data very simply using MATLAB's *polyfit* function, which can fit a polynomial function to a set of data. Since we wish want a linear function, we indicate via the third argument that only one variable is to be used. The first command below will regress values in a vector *y* on those in a vector *x,* giving a two-element vector *b* with the value of the slope coefficient followed by that of the intercept. The second command computes a vector of residuals (deviations of *y* values from the fitted amounts), and the third compute the R-squared value.

```
b = polyfit( x, y, 1 );
resids = y - ( b(2) + b(1)*x );
r2 = 1 - ( var(resids) / var(y) );
```

Statisticians will note that the R-squared value is not adjusted for degrees of freedom, a procedure not particularly germane in this context. In any event, the effect would not be noticeable since we have 100,000 data points.

It is also possible to fit a regression line to data points in a MATLAB plot using the Basic Fitting tool.

Additional Financial Considerations

We have now covered some fundamental components that will be used in subsequent chapters to analyze various retirement income strategies. But more can be done. We conclude this chapter with four major issues that are germane for many retirees but will not be included in our formal analyses.

First are issues associated with taxes. In the United States, income from capital gains is usually taxed at different rates than those used for income from dividends. Moreover, amounts received from certain bonds issued by states, local governments and some public agencies may also be taxed at preferential rates. Especially important is the fact that in the United States, some wealth may be held in accounts that provide tax exemption for the receipt of interest and dividends as well as proceeds from sales of securities, mutual funds and ETFs (but only while the funds remain in the accounts). Prominent examples are 401(k) and 403(b) retirement accounts, various types of Individual Retirement Accounts (IRAs) and other vehicles with funds intended for use in retirement. With many such accounts, income tax must be paid on amounts withdrawn to be either spent or invested in another fund without such favored tax treatment. The United States tax code also requires that certain minimum amounts be withdrawn from some of these tax-deferred accounts as the beneficiary ages, with income tax paid on the proceeds. We will analyze income strategies using such *required minimum withdrawals* in later chapter, but without taking into account any effects of differential taxation.

Another aspect that we will leave for future research is the impact of owner-occupied real estate, including primary and secondary residences. For many retirees, the value of home equity may be even larger than the amount of discretionary savings. While such real estate is likely to generate expenses for taxes, maintenance and possibly mortgage payments, it may also provide a source of needed funds, especially later in life. One may be able to borrow money by purchasing a new mortgage or increasing the amount of a current one, using the home as collateral. Another alternative is a *reverse mortgage*, which requires no monthly mortgage payments from the owners but grants the lender an option to purchase the house at a low price when the current occupants die or voluntarily vacate the premises.

A third aspect of retirement that we will not include is the possible need for *long-term care* for one or both retirees. Such care may be provided by family members, friends or by paid part-time or full-time caregivers in one's home. In the latter event, there may be substantial costs involved. Alternatively, one or both retirees may move to a retirement facility that provides such care, such as a community with facilities for *independent living*, *assisted living* and *skilled nursing*. Monthly fees at such communities can be substantial; many require a large initial payment as well. Some retirees are able to use equity in their home to offset at least some such costs. It may also be possible to use a reverse mortgage to help pay for home care. Or a move to a retirement community may be financed in whole or in part by selling one's home.

Some future expenses may be financed via the purchase of a l*ong-term care insurance policy,* covering some or all the increased costs incurred when one or both retirees is unable to independently perform a pre-specified number of *activities of daily living.* Such "*ADL*s" typically include eating, bathing, dressing, toileting, walking and maintaining continence. Each will be defined in a long-term care policy in excruciating detail, along with procedures for assessment of the severity of such problems. Long-term care policies differ in many respects, including the number of years covered, amounts to be paid, provider rating and cost. Estimates of the probability that a person will at some point require long term care, the likely duration of the need and the associated costs differ, but probabilities are not trivial and costs can be substantial. Nonetheless, only a minority of retirees currently insure against such risks. This may be explained in part by the availability of government assistance when such care is needed. In the United States, the Federal Medicare program does not cover many long-term care costs, but the Federal Medicaid program (or a variant offered in conjunction with a state government) will provide long-term residential care for those with verifiably few assets. However, the amount paid to a providing facility is likely to be relatively small and the conditions there often spartan, leading some retirees to exhibit what has been termed "Medicare Aversion".

A final aspect that we will not cover is the possibility of an explicit or implicit agreement between retirees and others (such as their children or other relatives) combining inheritance and payments for long-term care or routine costs that could be needed at advanced ages. Many retirees plan to leave some or all unspent wealth to family members in return for an understanding that if needed, some or all of the them will provide funds and/or help with needed care. In effect, mortality risk and uncovered health risks are pooled within a larger group. Of course such agreements are typically not binding and are subject to the criticism that they are "not worth the paper they are not written on", but these situations are not uncommon.

As we suggested in Chapter 4, one could include additional personal states covering the need for long-term care in the client personal state matrix, given sufficient actuarial data to make reasonable probabilistic projections. One could also model the probability distribution of the value of particular real estate holdings, including correlations of such values with the overall market portfolio. And, given sufficient patience, the effects of differential taxation could be included, and even probabilities of possible future changes in tax codes.

To summarize: there is clearly more work to be done analyzing the key elements of retirement income scenarios. But we leave much of this for others, proceeding to analyze the attributes of retirement income strategies in a relatively simple world.