Retirement Income Analysis
with scenario matrices

William F. Sharpe

15. Lockboxes

Theory and Practice

This chapter is about investment strategies that use vehicles that we will call *Lockboxes,* as described in my working paper "*Lockbox Separation*" in June 2007 (available at *www.stanford.edu/~wfsharpe*) and a joint paper with Jason Scott and John Watson. "*Efficient Retirement Financial Strategies,*" written in 2007 and published in 2008 in John Americks and Olivia Mitchell's book "*Recalibrating Retirement Spending and Saving*".

Initially, the discussion will be theoretical (antonym: realistic). This is a practice often engaged in by economists to simplify analysis and focus on key aspects of a problem. The reader's indulgence is requested in the hope that the key ideas will lead to useful and practical investment products and/or services.

Lockbox Contents

To begin, let's focus on the provision of income in a specific future year – for example, year 10 (9 years hence). To keep things simple, let's also assume that both Bob and Sue will be alive at that time. Today we create *Lockbox10* to provide their income in year 10. And today we put into the box, chosen amounts of some or all of three types of investments:

1. Zero-coupon TIPS maturing in 9 years
2. World Bond/Stock Mutual Fund or ETF Shares to be sold in 9 years
3. m-Shares maturing in 9 years

The box is to be sealed after the contents are put in, then opened at maturity (here, 9 years).

You were warned that this would be theoretical. Let's see why.

First, at present every maturity of TIPS securities provides coupon payments (although the coupons are relative low for some newer issues due to the exceedingly low interest rates in recent years). Second, there may be no outstanding TIPS issues with maturities for some future years. Third, at the time of this writing there is no single World Bond/Stock mutual fund or ETF (although one can be simulated using the procedures described in Chapter 7). And fourth, there are currently (in 2017) no m-shares *per se*, in the sense described briefly in Chapter 9. But our lockboxes are in part an aspirational concept, so please keep reading.

m-shares

An m-share is a security that promises to pay a real amount per share at a single pre-specified *maturity date,* with the amount paid being a *non-decreasing function* of the *cumulative real return* on the *market portfolio* from the present to the maturity date. Here, as throughout this book, the cumulative real return on the market portfolio at a given date equals the real value at a future date of $1 invested in the market today, so the cumulative return cannot be negative as long as the securities in the market portfolio have limited liability. As a practical matter, we assume that the World Bond/Stock Mutual fund is a sufficient proxy for the market portfolio.

As discussed in Chapter 9, a financial service firm could create any desired type of m-share by purchasing a portfolio of TIPS and the market portfolio and issuing two classes of shares. The first class would make the payments required for the desired m-share; the other class would make payments from the assets remaining after the first class was paid. While not absolutely necessary, it is preferable for both classes to be m-shares, with payments that are non-decreasing functions of the cumulative return on the market portfolio.

Below is an illustration of the basic approach. The green curve shows a desired payout for a 10-year security that we will call m-share A. As discussed in Chapter 9, this is equivalent to (1) purchasing market shares, (2) selling an option for someone to call the shares at a higher price, and (3) purchasing an option to allow the holder to sell shares at a lower price. As we know, in the option trade such an approach is sometimes called an *Egyptian* strategy.

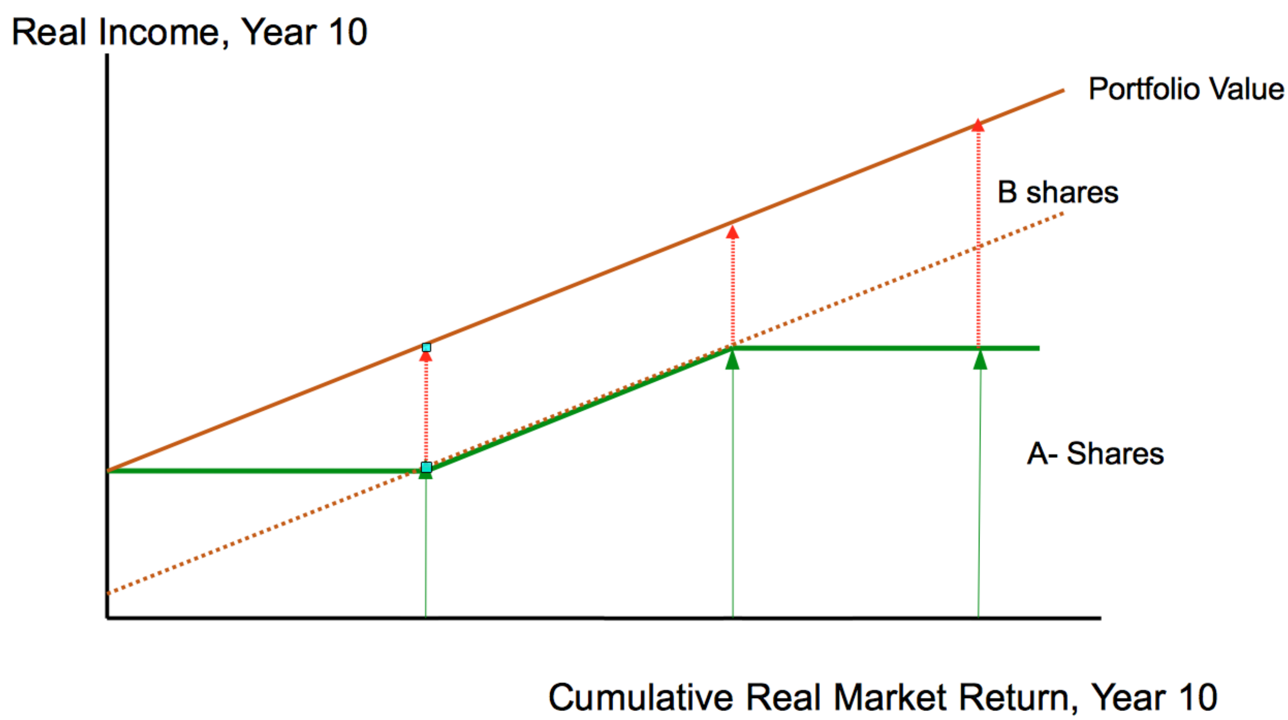

Now, find the steepest slope along the green curve. Here it is equal to the slope shown by the dotted red line. Next, move this line up until it lies on or above the green line for every value on the x-axis. Here we use the lowest such line, shown by the solid red line, but we could have used a higher parallel line. We then purchase a combination of the risk-free asset and the market portfolio that will provide the real incomes shown by the solid red line. Then we issue two classes of shares: A and B. At the maturity date, we pay out the entire value of the fund to the two share classes, with the amount shown by the green curve to class A and the remainder to class B. Note that both classes are m-shares since the real income of each is a non-decreasing function of the cumulative real market return at maturity. As we have indicated, Class A provides payments equal to that of an Egyptian strategy: going from left to right the curve is flat, then up, then flat *(fuf)*. And inspection shows that Class B provides payments equal to that which, as indicated in Chapter 9, is sometimes called a *Travolta* strategy. The amount paid to Class B shares, if shown separately would plot on a curve that, going from left to right would go up, then be flat, and then go up again (*ufu*). Note that in order to qualify as a non-decreasing *function* of the cumulative market return, an m-share's curve cannot be vertical or downward-sloping.

To generalize: A financial institution can create any desired type of m-share by following this approach. The result can provide one class of m-shares with the desired payout structure, and another class with a complementary structure. Absent outright fraud, there should be no default risk. And, given sufficient competition, overall expenses (fees, etc.) should be very low.

This example illustrates another important point. For every investor who wishes to have a payout that is a non-linear function of the return on the overall market, there must be one or more others willing to accept a payout that is a complementary function of the market return. For example, investors who want Egyptians need others willing to accept Travoltas, and vice-versa. The prices of the two share classes will need to adjust as needed to clear the markets and, given any reasonable sort of equilibrium, the values of the classes should be close to the value of the underlying pool of TIPS and market portfolio shares used to create the m-shares.

The term "*non-decreasing function*" is cumbersome but essential. Consider a graph such as the one above, with the terminal value of the market portfolio at a given time on the x-axis and the terminal value of the m-share at that time on the y-axis. It must be possible to graph the payments made by the m-share by putting a pen (or stylus) at the origin, then moving it to the right and either horizontally or upward, but never vertically or to the left, all the while not picking the pen up until reaching the right side of the graph. This is a necessary and sufficient condition for the value of the m-share to be a non-decreasing function of the terminal market value.

Note that the market portfolio meets our definition of an m-share, as does a riskless real asset. In fact, we could have defined our lockbox as simply a box holding an m-share. But we choose to differentiate the three possible investments, restricting the term "m-share" to describe an instrument with payments that plot as a *non-linear* and non-decreasing function of the cumulative return on the market portfolio.

Cost Efficiency

In Chapter 8 we showed that the least-cost way to obtain any given set of possible incomes in a year is to allocate the payments across scenarios so the amount of income is a non-increasing function of price per chance. We now show that the income produced by any investment in our type of lockbox is 100% cost efficient in this sense.

Consider the situation shown below. A couple has decided that $60,000 per year from savings, plus income from Social Security will provide a satisfactory standard of living in year 10. Accordingly, they invest in an m-share that will pay the amounts shown below, depending on the cumulative market return in the next nine years.

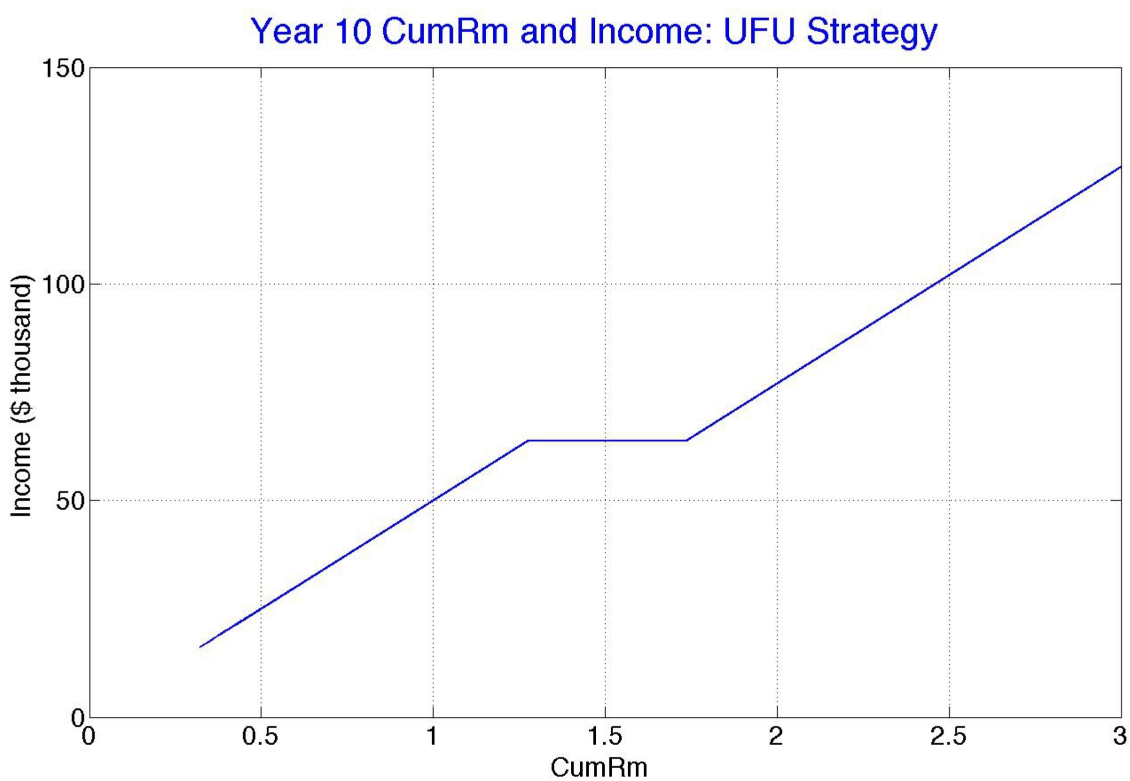

The terms of this m-share have been constructed so that there is a 33.3% chance that income will be below $60,000, a 33.4% chance that it will equal $60,000 and a 33.33% chance that it will exceed $60,000. In financial jargon, it is a Travolta.

This clearly meets the requirement for an m-share: income is a non-decreasing function of the cumulative return on the market. And, of course, each of the other two instruments allowed in our lockbox – investments in the market portfolio or TIPS would also provide income that is a non-decreasing function of the cumulative return on the market..

Recall the relationship between price per chance (PPC) and the cumulative return on the market portfolio: PPC is a decreasing function the return on the market, as shown below for year 10:

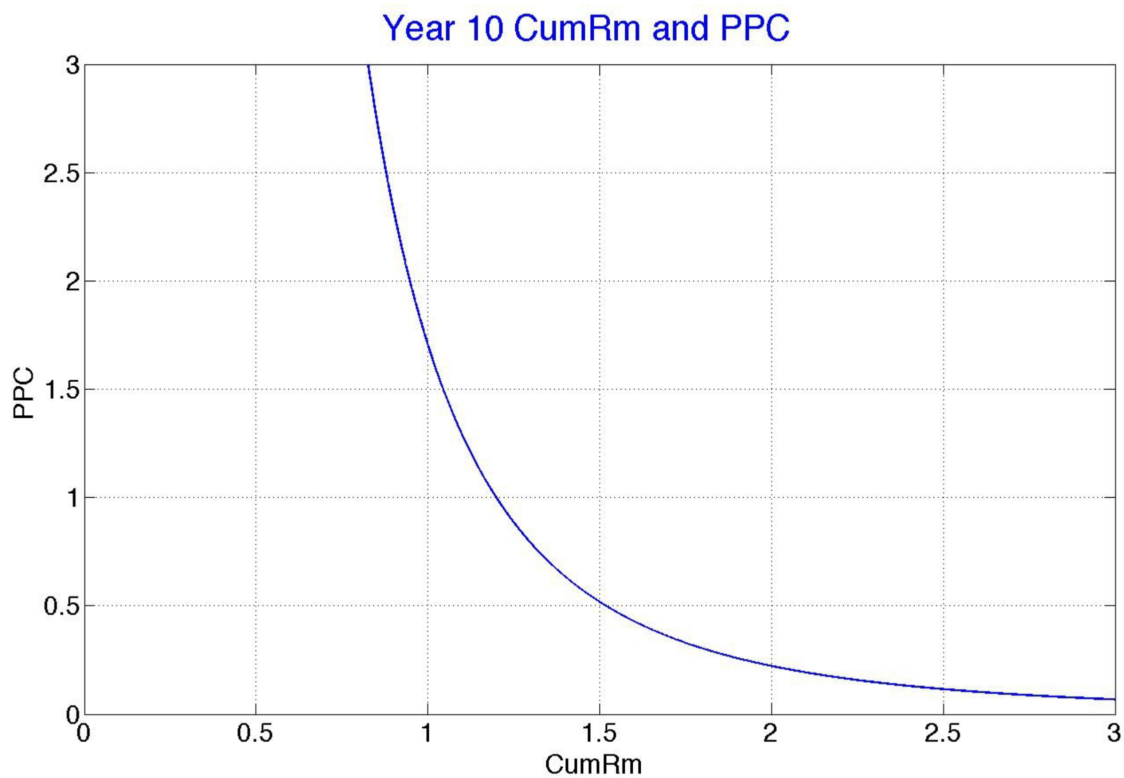

Combining the relationships in the two previous graphs gives the following:

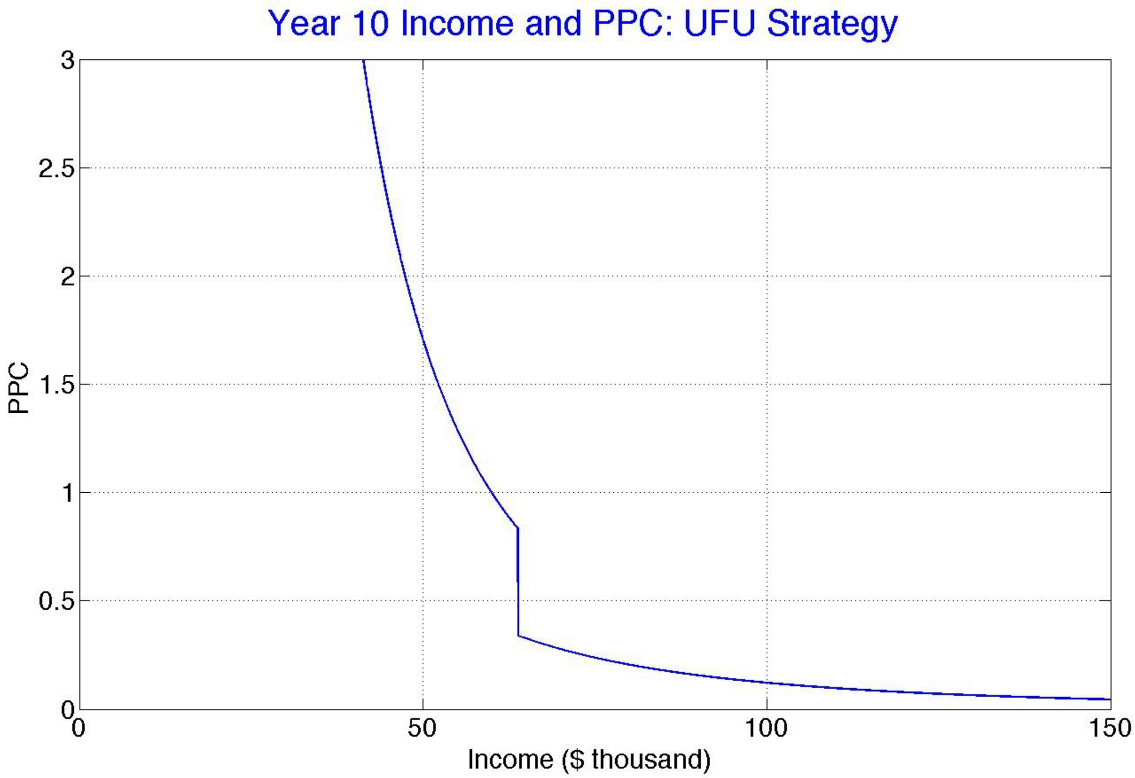

Income is indeed a non-increasing function of PPC. And there is thus no cheaper way to obtain the chosen distribution of income. Thus the Travolta strategy's cost-efficiency is 100%.

This result is more general:

1. The investments in *any* of our lockboxes will be cost-efficient, and
2. *any* distribution of incomes in a year can be obtained at lowest cost using such a lockbox strategy.

Our type of lockbox strategy or the equivalent is thus a *necessary and sufficient* condition for 100% cost-efficiency.

One note is in order before we continue. We have not considered *dynamic strategies* that adjust holdings of a risky portfolio and a riskless asset as values change, with the hope of obtaining a terminal value that is close to some pre-specified function of market return. The omission is intentional. In frictionless markets that allow frequent trades at prices that change by tiny increments, such strategies could provide cost-efficient results. But in actual markets, this is unlikely – the results could at best approximate a desired function and substantial costs would be incurred for frequent transactions.

Utility

There is more. Chapter 9 showed that a person wishing to maximize the expected utility of income in a given year should choose a strategy for which the utility of income in each possible future state is equal to a state's price per chance times some positive constant plus some other positive constant. Thus a graph such as the one above can be taken to represent the relationship with the recipients' marginal utility of income in the year in question. It thus reveals important information about preferences.

A particularly interesting aspect of this example is the range of PPC values for which the recipients have chosen a constant amount of real income. The largest PPC in this range is 2.46 times the smallest. This implies that the marginal decrease in utility from a slight decrease in income from $60,000 is almost 2.5 times as large as the increase in utility from a slight increase in income from that level. Using terminology from behavioral economics, the pain from a small decrease in income from the *reference poin*t of $60,000 is 2.5 times as great as is the pleasure from an increase of a similar magnitude. The recipients' underlying utility curve thus has a *kink* at the reference point, leading to a reluctance to accept lower incomes unless the cost is very large and higher incomes unless the cost is considerably lower.

This sort of behavior is consistent with key aspects of the approach presented by Amos Tversky and Daniel Kahneman in their seminal 1979 paper "*Prospect Theory: An Analysis of Decision Under Risk*". Moreover, the ratio of implied marginal utilities at the *reference point* is close to the magnitudes frequently implied by choices made by subjects in empirical studies. Such experiments tend to suggest that for many people the *displeasure* from a small *loss* relative to the reference point appears to be two to three times as large as the *pleasure* from a small *gain* from that point.

Tragically, Amos Tversky died in 1996 at the age of 59. In 2002, Daniel Kahneman received the Nobel Prize in Economics for their joint work (the Prize is not awarded posthumously). Perhaps we should follow suit and call a strategy such the one we have been analyzing a *Kahneman* instead of a *Travolta*. Rather than taking a position either way, we will refer to such an approach by the key characteristics of its plot: moving from left to right, it goes up, then remains flat, then goes up again. Thus, up-flat-up or UFU – pronunciation: *oo-foo*. The plot for the complementary strategy, heretofore called an Egyptian, is flat, then goes up and then is flat again, hence: FUF (rhymes with *muff*).

While an UFU strategy may be preferred by some investors, it of course does not offer something for nothing. The figure below contrasts it with two equal-cost alternatives: 100% investment in the market portfolio and 100% investment in the riskless real asset. It may seem strange that the blue curve is only slightly above the red for lower market returns and well below it for larger returns. But this reflects the fact that money in low future market return states is considerably more expensive than money in high market return states. In competitive capital markets, there are no free lunches.

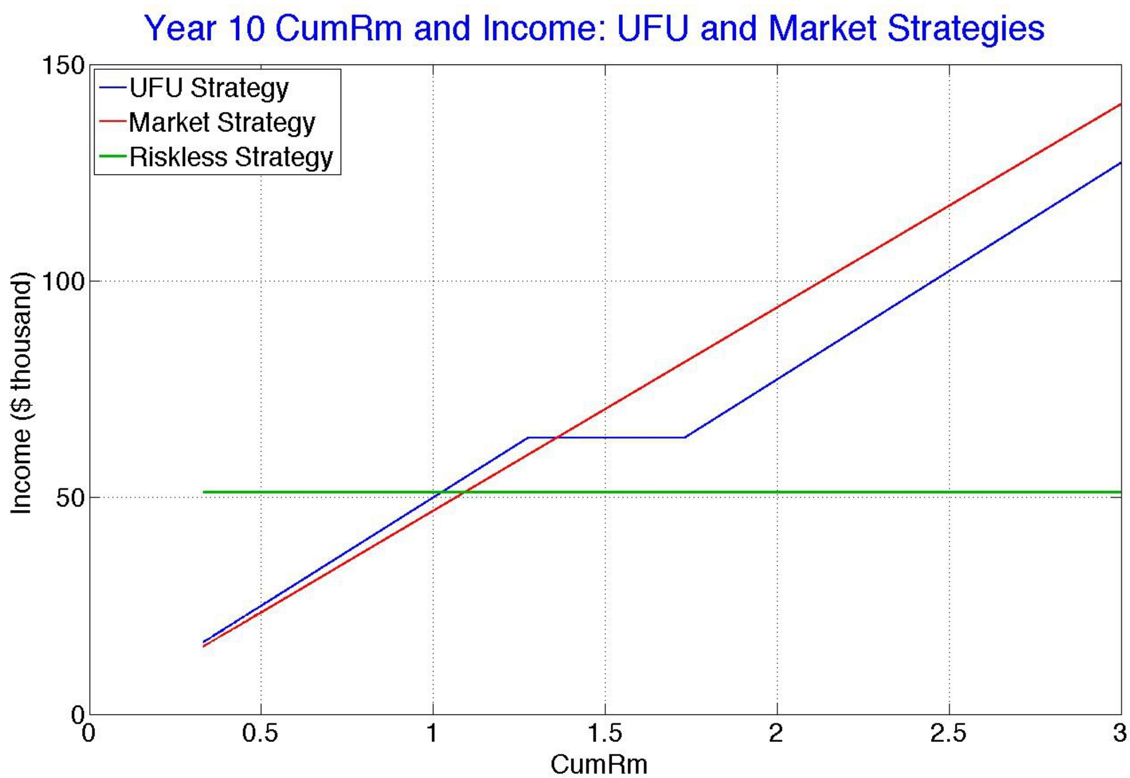

Income Distributions and m-share Terms

In our ideal (theoretical) world, an investor can obtain *any* set of income probability distributions for future years that he or she can afford. In this section we will see how this could be done.

Consider the following situation. An investor has created a lockbox that will mature in year 2, with $1,000 invested entirely in the market portfolio fund. The probability distribution for its value at maturity is shown below.

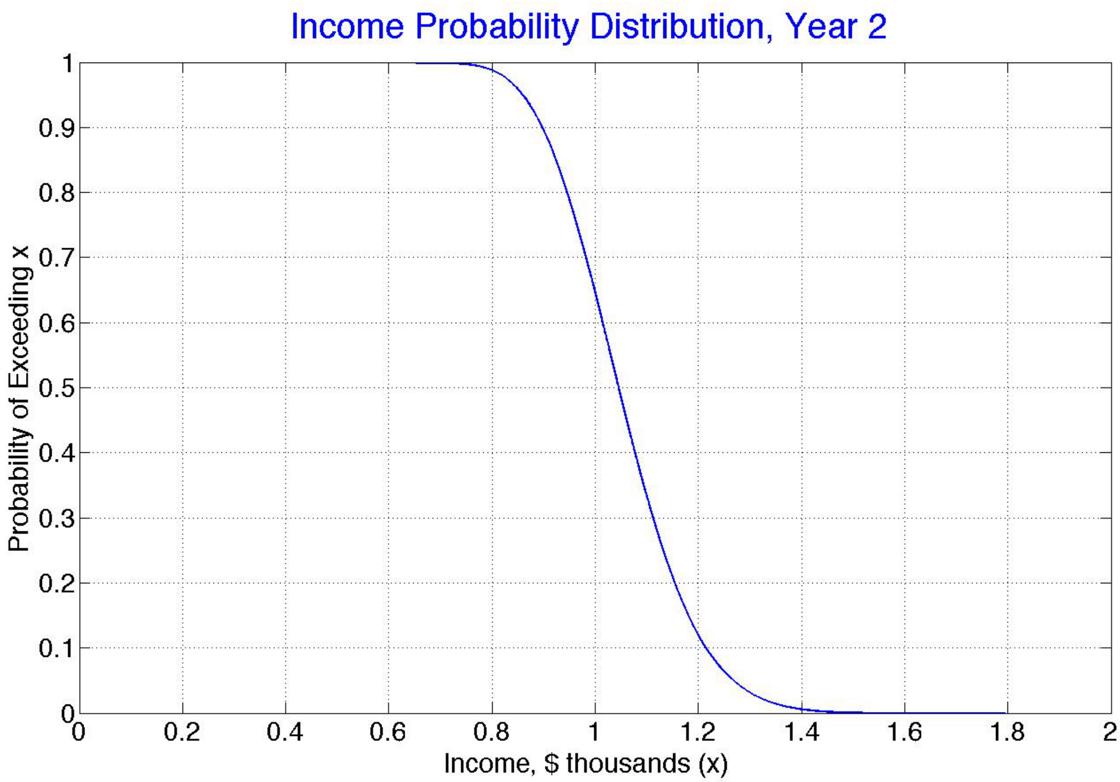

The investor would like to have the same probability distribution of income (as seen from today) for every other future year. We will use the lockbox for year 10 as an example. The problem is that the cumulative returns on the market portfolio through year 10 are not distributed in the same manner as those through year 2. But there is a way that an m-share class could be created to provide similar probability distributions of returns in the two years.

We start with the creation and processing of a market data structure. From this we extract from the *market.cumRmsM* matrix column 2 with 100,000 possible cumulative returns through year 2 and column 10 with 100,000 possible cumulative returns through year 10. The curve in the graph below shows a cross-plot of the two sets of returns, with each sorted from lowest to highest. Now, assume that this curve shows the terms of an m-share. In year 10 the issuer would plot the realized cumulative return on the market portfolio on the horizontal axis, then pay an amount equal to the corresponding amount on the vertical axis. If future returns are drawn from the 100,000 scenarios used for the construction of the m-share, it would offer the same *ex ante* distribution of incomes in year 10 as did the market investment in year 2.

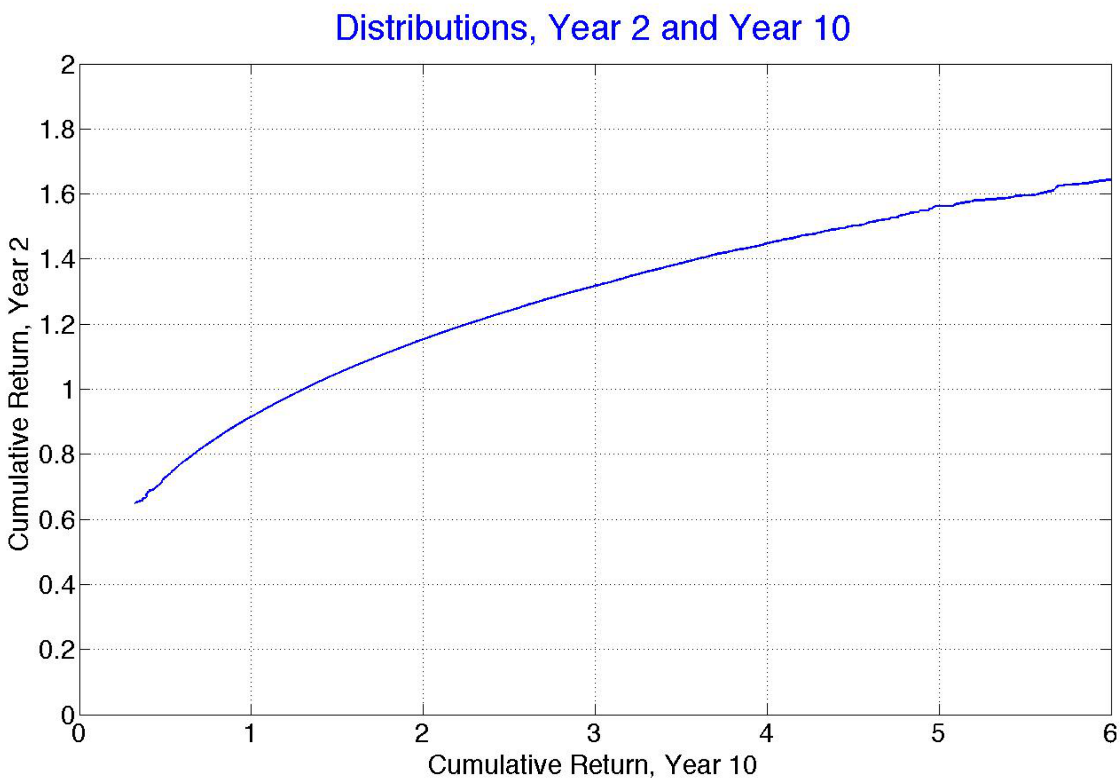

But this is a big if. At the very least, we should examine the possible results provided by such an m-share using a newly generated matrix of market portfolio returns. And the results are likely to differ, if only slightly, from those on the curve showing its terms. For this reason it might be best to generate the curve using a larger number of scenarios and to smooth it somewhat, especially at the values corresponding to very large and very small cumulative market returns.

These caveats aside, viewed from today, such an m-share could provide a probability distribution of possible cumulative returns 10 years hence very similar to the probability distribution of possible cumulative returns 2 years hence.

More generally one could, in principle, design an m-share with a distribution of terminal values approximately equal to any type desired by enough investors to warrant the effort.

Linear Approximations of Income Distributions

At present there are few, if any, financial instruments can provide payments that are non-linear functions of the return on a broad bond/stock market portfolio over a period of many years. It is possible that eventually there might be sufficient demand for an investment firm to provide such securities using default-free and low-cost vehicles such as m-shares. But at present such low-cost and simple instruments are unavailable. An alternative is to create a lockbox with only TIPS and/or the market portfolio to provide a linear approximation of a desired distribution of income in a future year.

One way to do this is shown below.

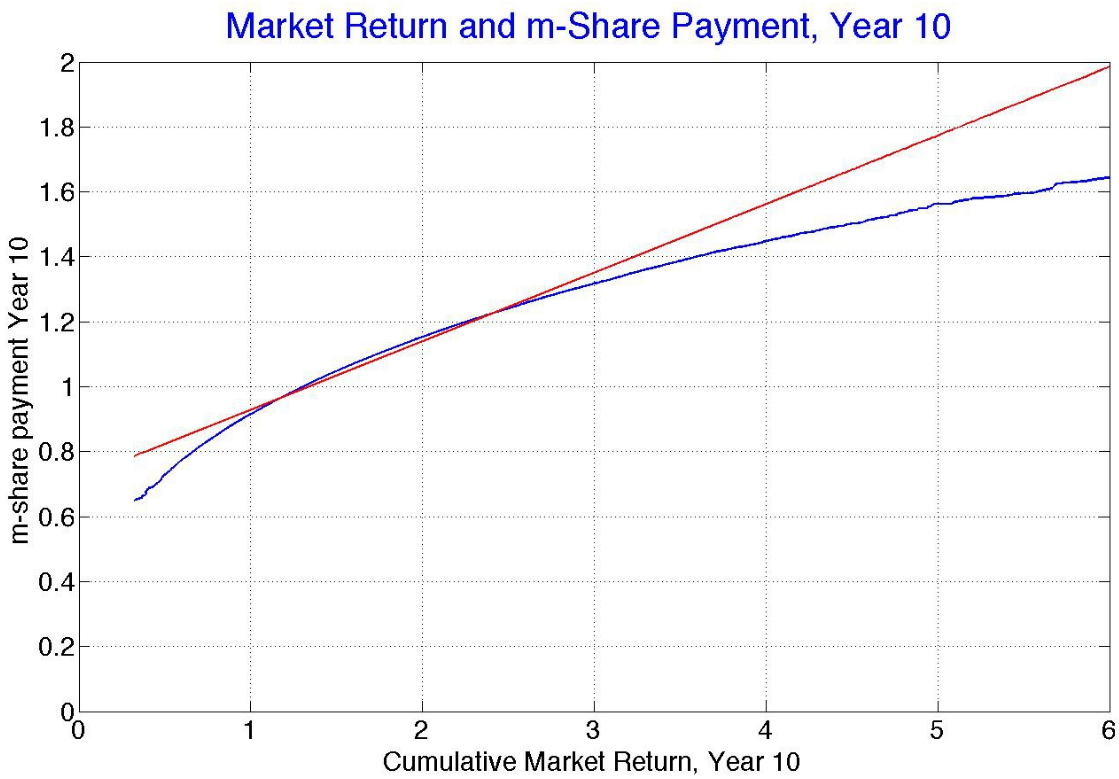

The blue curve is the same one shown earlier, but the y-axis has been labeled as the m-share payment in year 10. In addition, we have fitted a line to the points on the blue curve using our standard commands for least-squares regression:

```
xvals = [ ones(length(x), 1) x ];
b = xvals \ y;
yFitted = b(1) + b(2)*x;
```

Here, x is the vector of the x-values for the blue curve and y is the vector of the corresponding y-values on the curve.

It may seem surprising that the red line seems to lie above the blue curve more than it lies below it. Why? Because there are fewer scenarios with cumulative market returns below 1.0 or above 3.0 than there are between 1.0 and 3.0, and each scenario is given equal weight when fitting the regression line, which is designed to minimize the sum of the squared deviations of the fitted (red) points from the original (blue) points.

In this case, the the value of b(1) is 0.7162 and that of b(2) is 0.2124. These have direct economic interpretations. We can approximate the distribution of cumulative market returns obtained with the investment of $1 in the market portfolio held for 1 year (until year 2) with an investment in the risk-free asset and the market portfolio held for 9 years (until year 10). Note that the intercept indicates the ending value if the cumulative return on the market portfolio is zero (that is, nothing is left from investing in the market). Thus the return on the risk-free holding is $0.7162. But we know that this asset has a real return of 1% per year. Thus the initial amount invested in the risk-free asset must equal 0.7162/ (1.01^9), or $0.6548.

Note also that if the cumulative return on the market is 1.0, the value of the m-share will equal 0.7162+0.2124*1. The difference between the ending value if the cumulative return on the market is 1.0 and the value if it is 0 is thus 0.2142. Therefore the amount invested in the market initially must be 0.2124. More generally:

$$\text{Initial investment in the risk-free asset} = \frac{b_1}{(1+rf)^t}$$

$$\text{Initial investment in the market portfolio} = b_2$$

In this case, the sum of the two amounts invested = 0.8672, so only 86.72% as much needs to be invested in lockbox10 as in lockbox2 in order to have roughly similar probability distributions of income in the two years.

Here are the actual distributions of income in years 2 and 10 for our two lockboxes, assuming the cost of the first is $1 and that of the second is $0.8672. As expected, they are similar, but not exactly the same.

AMDnLockboxes

It is useful to generalize the prior example. First, we can consider a set of lockboxes, maturing at the beginning of years 1, 2 and so on. As in the case shown for year 10, each can be designed to provide a distribution of values similar to that of a lockbox with $1 invested at present in the market portfolio and maturing a year hence at the beginning of year 2. We call such a set: *AMD2 lockboxes*.

By extension, we will also consider strategies designed to approximate returns obtained by holding the market portfolio for more years – hence *AMD3*, *AMD4* and so on. More generally, an *AMDn strategy* is designed to provide for each year after year *n*, an income distribution approximately equal to the income distribution obtained by holding the market portfolio until the beginning of year *n.* Since we do not allow borrowing at the TIPS rate of interest, for cases in which *n* is greater than 2, we assume that the market portfolio and/or TIPS are held in each lockbox maturing before or at year *n*.

Since AMDn lockboxes can be used for annuities or for strategies that do not involve annuitization, it is convenient to be able to create and process AMDnLockbox data structures that could be utilized for either purposes. Our goal is to create a generic version of such a set of lockboxes, with the total value invested in the first lockbox equal to 1.0. Subsequent functions can then scale the values in each of the boxes, as needed.

Here is the *AMDnLockboxes_create* function:

```
function AMDnLockboxes = AMDnLockboxes_create( );
  % creates an AMDn lockboxes data structure

    % year of cumulative market return distribution to approximate (n)
    % note: n must be greater or equal to 2
      AMDnLockboxes.cumRmDistributionYear = 2;

   % lockbox proportions (computed by AMDnLockboxes_process)
      AMDnLockboxes.proportions = [  ];

    % show lockbox contents (y or n)
      AMDnLockboxes.showProportions = 'n';

end
```

The first parameter indicates the desired year's distribution to be approximated. The second provides a data element that will contain the proportions after *AMDnLockboxes_process* is run. And the last parameter indicates whether or not a graph of the results is to be shown.

The *AMDnLockboxes_process* function has two main sections. The first does the calculations, the second provides a graph if one is requested.

Here are the initial statements:

```
function AMDnLockboxes = AMDnLockboxes_process(AMDnLockboxes, market, client);

 % get number of years of returns
   [nscen nyrs] = size( market.cumRmsM );

 % get n
   n = AMDnLockboxes.cumRmDistributionYear;
   if n < 2 ; n = 2; end;
   if n > nyrs; n = nyrs; end;

 % set lockbox proportions for initial years to investment in the market portfolio
   xfs = zeros( 1, n-1 );
   xms = ones ( 1, n-1 );
 % create matrix of proportions
   xs = [ xfs; xms ];
```

This section creates the initial matrix of values for each lockbox maturity year, with the amounts for TIPS holdings in the initial row and the amounts for the market in the second row. While the holdings for year 1 (which will be spent immediately) could be any combination of TIPS and the market portfolio that sums to 1.0, we arbitrarily choose the market portfolio. We also use it for any subsequent maturity year prior to the year *n-1*. And, just to be safe, we insure that the values of *n* are within allowable bounds.

These tasks complete, the function next computes the required contents for each of the subsequent lockboxes. For each one, we create a vector *x* of sorted cumulative market returns for the base year *n* and a vector *y* of sorted cumulative market returns for the year in question. Then, as in the earlier example, we use regression analysis to find the parameters of the least-squares linear relationship between the two sets of values. The next statements compute the amounts to be invested in the riskfree asset and the market portfolio, using our prior results, then add them to the matrix *xs*.

```
% do regressions to compute contents of remaining lockboxes
  for yr = n: nyrs
    % sort cumulative returns
       x = sort( market.cumRmsM( : , yr ), 'ascend' );
       y = sort( market.cumRmsM( : , n ), 'ascend' );
    % regress y values on x values
    %     y = b(1) + b(2)*x
       xvals = [ ones(length(x), 1)  x ];
       b = xvals \ y;
    % compute lockbox contents
      xf =  b(1) / mean( market.cumRfsM ( : , yr ) );
      xm =  b(2);
    % add to xs matrix
      xs = [ xs  [ xf ;  xm ]  ];
  end % for yr = n: nyrs
```

When all the lockbox contents have been computed, the resulting matrix is placed in the element *proportions* of the *AMDnLockboxes* data structure so that it can be used by other functions to produce incomes:

```
% add lockbox proportions to  AMDnLockboxes
  AMDnLockboxes.proportions = xs;
```

The remaining statements of the function produce a graph if desired, then end the function:

```
  % plot contents if requested
    if lower( AMDnLockboxes.showProportions ) ==  'y'
      fig = figure;
      x = 1: 1: size(xs,2);
      bar( x, xs' , 'stacked' ); grid;
      set( gca, 'FontSize' , 30 );
      ss = client.figurePosition;
      set( gcf,  'Position' , ss );
      set( gcf, 'Color' , [1 1 1]  );
      xlabel(  'Lockbox Maturity Year ' , 'fontsize' , 30 );
      ylabel(  'Amount Invested at Inception    ' , 'fontsize' , 30 );
      legend( 'TIPS ' , 'Market ' );
     ax = axis; ax(1) = 0; ax(2) = nyrs+1; ax(3) = 0; ax(4) = 1; axis(ax);
      t = [ 'Lockbox Proportions for approximating Market Distribution in year ' num2str(n) ];
     title( t, 'Fontsize' , 40, 'Color', 'b' );
     beep;  pause;
    end; %if lower( AMDnLockboxes.showProportions) = 'y'

end % function
```

And here is the graph produced by the function for a case with $n = 2$:

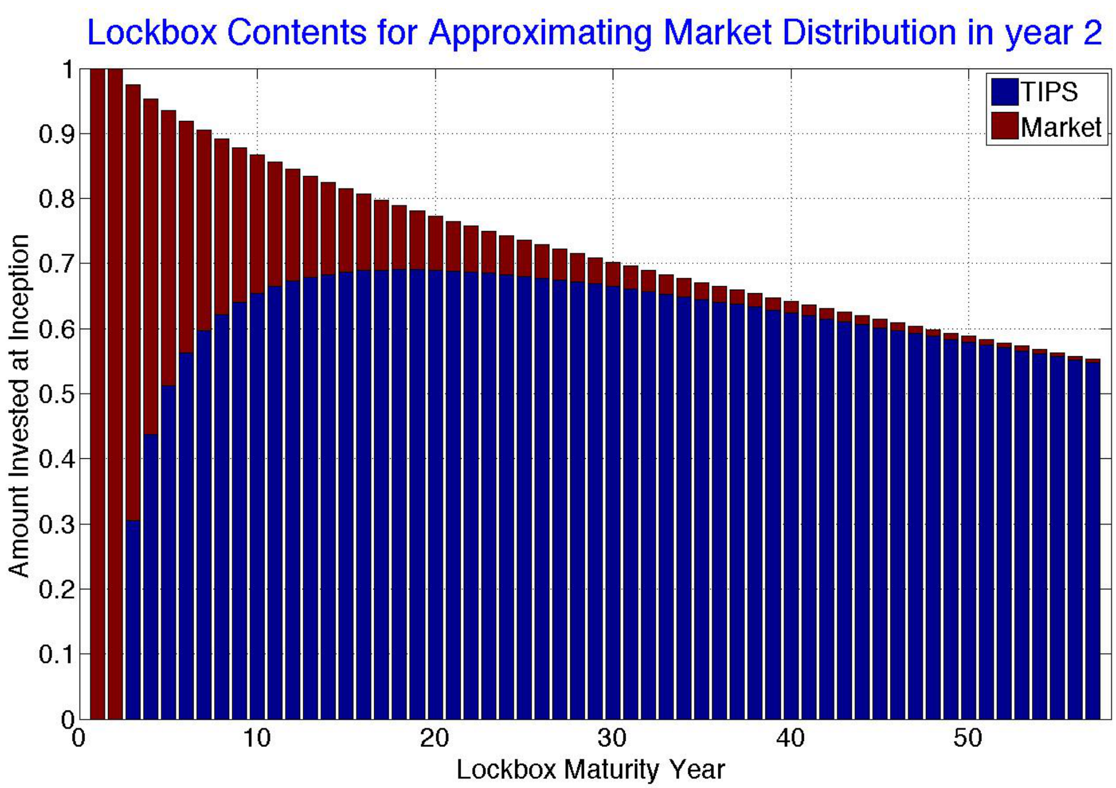

Before proceeding, we need to consider the issue of sample bias. In principle, we should not use a matrix of possible cumulative market returns to construct our lockboxes, then analyze their performance using the same matrix of cumulative market returns. Any such matrix should be considered a sample of scenarios from the larger population of a great many possible scenarios. Ideally we would construct our lockboxes analytically using the formulas for the underlying distributions. Alternatively, we could employ a simulation using the best possible sample of the population of scenarios, but this could require a huge matrix with millions or billions of rows. That said, we could at least construct the lockboxes with one sample of scenarios, then apply the results using another. For example, consider a case in which we create our usual market data structure using the statement:

```
market = market_process( market, client );
```

then create a second market structure using the statement:

```
market2 = market_process( market, client );
```

This will use the same parameters for the risk-free real rate of return, market expected risk and return, etc. as in the earlier statement, but produce a different matrix of market cumulative returns. We would then create our lockbox contents using this new matrix:

```
AMDnLockboxes = AMDnLockboxes_process( AMDnLockboxes,  market2 );
```

The data structure *market2* can then be discarded, with the original market structure used for subsequent computations. This will yield results that are still subject to error but are unbiased.

While this is a better alternative, it may be relatively harmless to simply use the same market data structure to create lockboxes and then to determine their performance. The following two graphs show the results of an experiment in which 100 different market structures with our parameters were used to compute lockbox contents. The two leftmost dots reflect values that are all the same. For each subsequent year the results vary, as can be seen from the slightly elongated vertical plots, but the variations are very small indeed. Here, as with the survival probabilities, sampling error may be a minor concern.

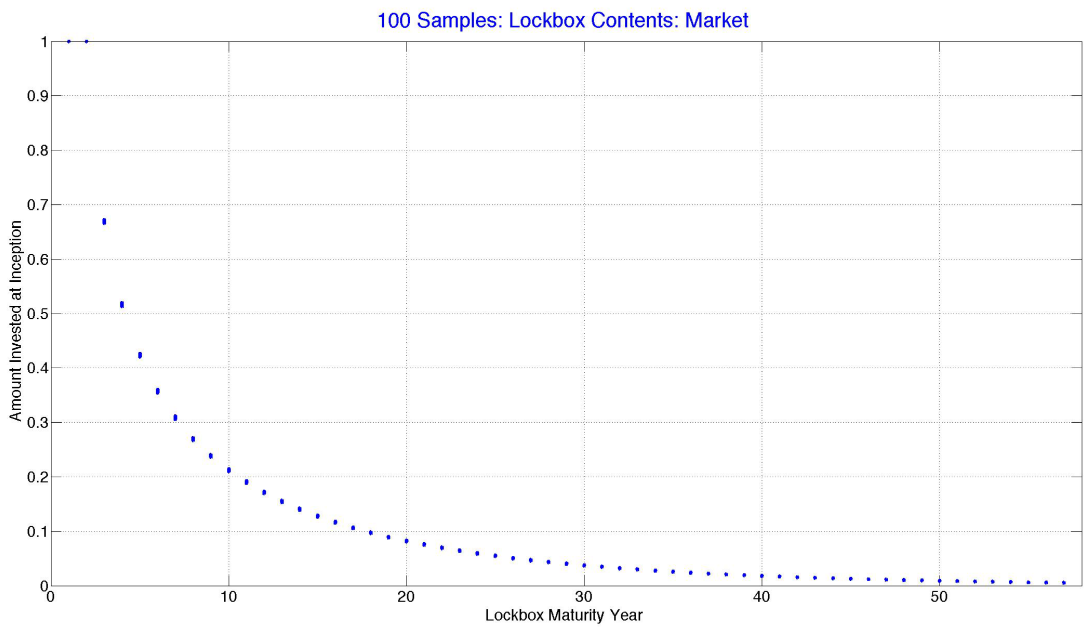

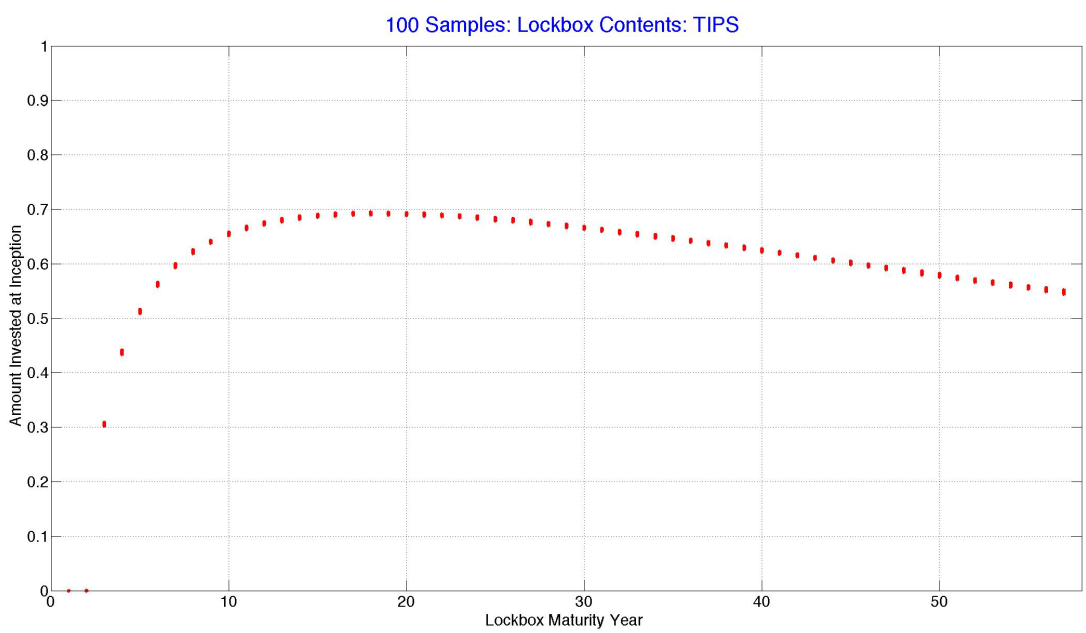

Returning to our main theme, consider a case with a single recipient named Angela. We assume that only Angela is alive (personal state 1) and that she will live for precisely 30 years. (Remember, this is still theoretical). She has enough money to invest $40,000 in lockbox 1 and proportionate amounts given by the lockbox contents shown in the bar chart above.

Here are the probability distributions of the values of the lockboxes at maturity, shown by the last version of the animated graph produced by setting the *analysis.plotIncomeDistributions* data element to 'y':

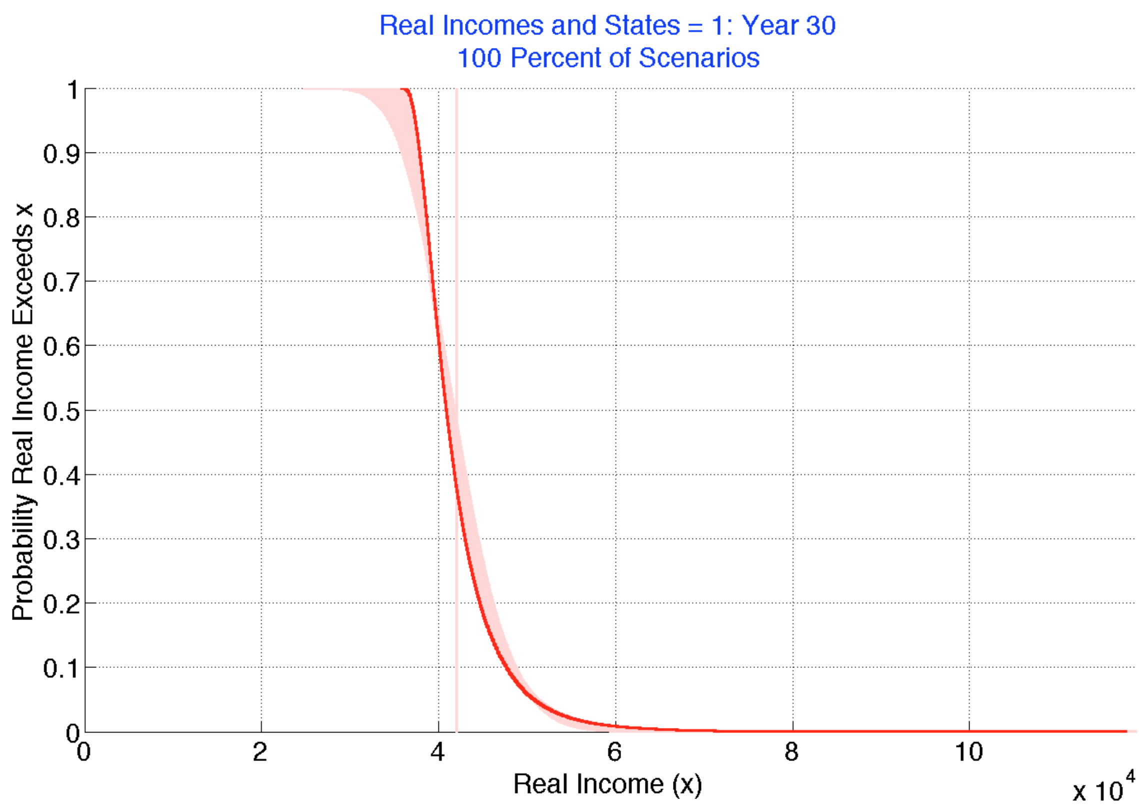

The probability distribution for year 2 has the lowest real income for high probabilities and very low probabilities, and it has the highest real income for mid-range probabilities. The distribution for year 30 has the highest real income for high probabilities, the lowest for mid-range probabilities and the highest for very low probabilities. Plots for other years fall neatly between these two. But the key result is that the distributions are, as desired, very similar.

The figure below shows the last version of the animated graph obtained by setting *analysis.plotYOYIncomes* to 'y':

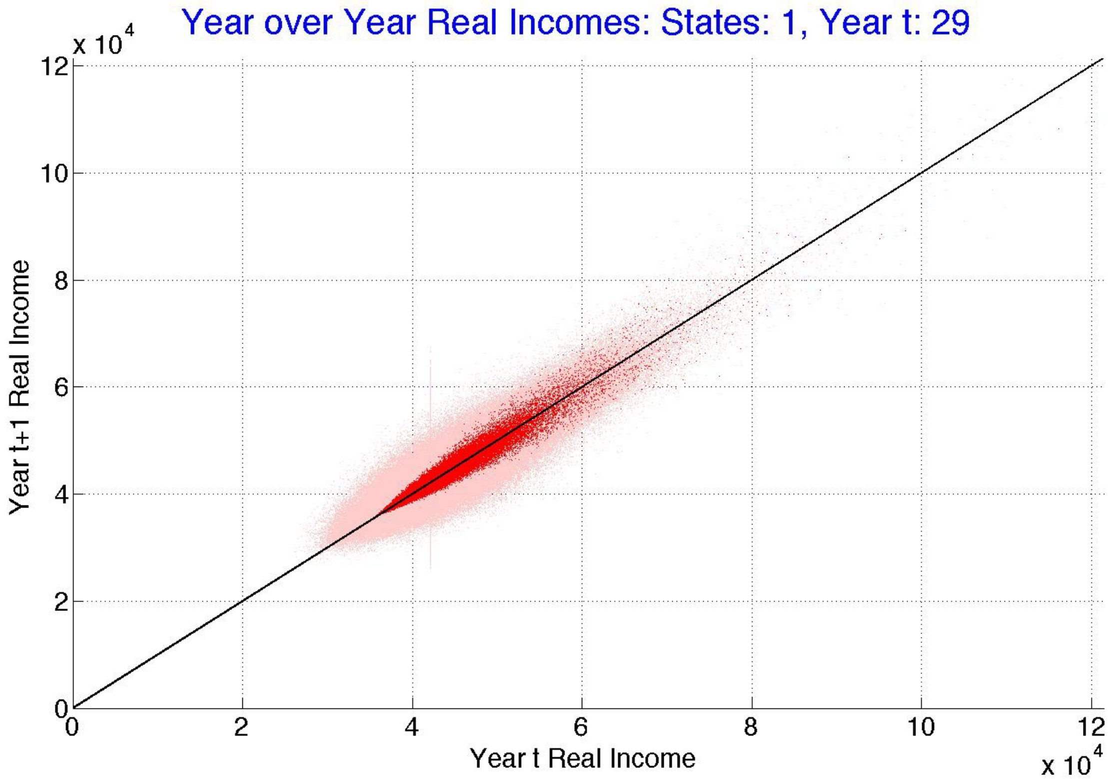

There can be substantial variation from year to year in the early years, but less in the later ones. Moreover, the variation tends to be smaller following years with either relatively low incomes or those with relatively high incomes.

Perhaps surprisingly, the relationship between real income and price per chance across scenarios (possible future states of the world) varies substantially from year to year, as shown by the following graph obtained by setting *analysis.plotPPCSandIncomes* to 'y' and, in order to obtain a plot with logarithmic values on both axes, setting *analysis.plotPPCSandIncomesSemilog* to 'n':

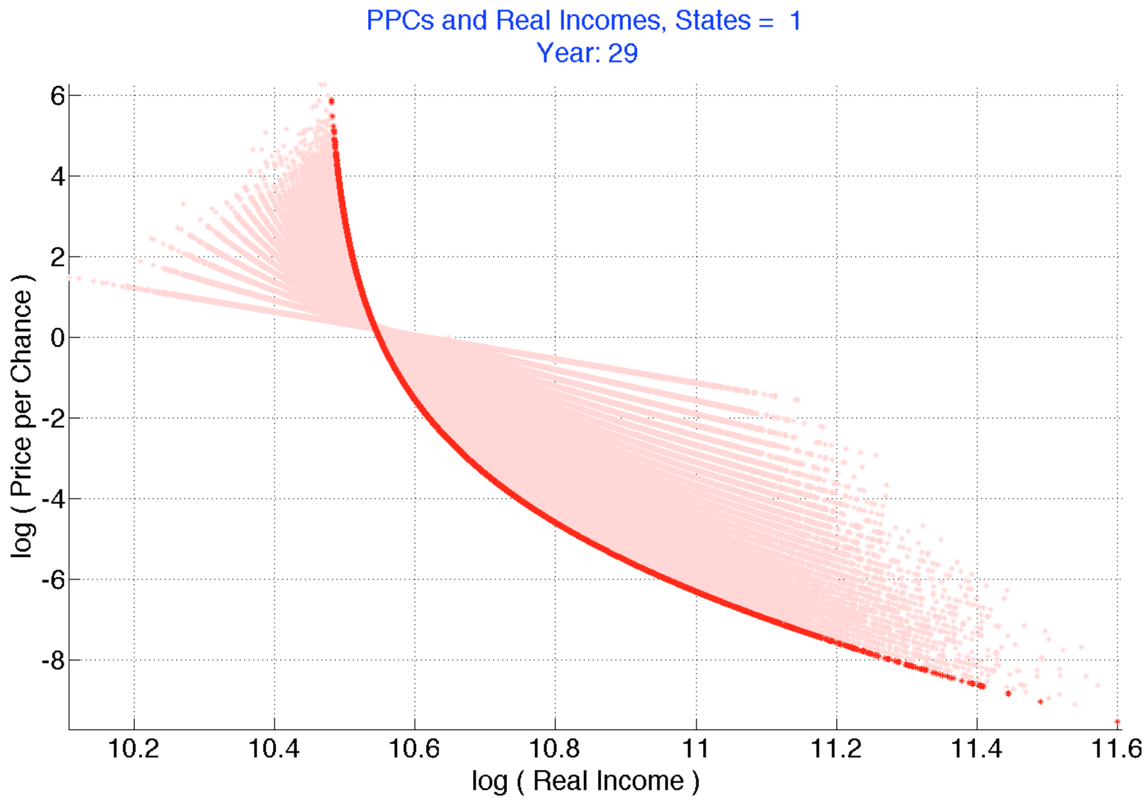

As with the previous graphs, this is the view near the end of the animation. The dark red curve shows the relationship for incomes in year 29. The flattest and shortest curve reflects the relationship for incomes in year 2. The remaining years fall between, covering larger income ranges, with greater slopes at the point at which the logarithm of PPC is 0 (and thus PPC = 1).

This may seem paradoxical. We assumed that Angela wanted similar probability distributions, viewed from today, for income in each future year. Yet her implied marginal utility functions for those years differ substantially. And this is not due in any way to mortality, since we have assumed that she is guaranteed to be alive through year 30.

Why? The formal answer is straightforward. The range of costs (price per chance) for income is wider for years farther in the future. Yet Angela has chosen nearly the same range of incomes. She takes differences in costs (PPCs) into account, but changes her planned spending less to respond to differences in costs in later than in earlier years.

Formally, only the curve for year 2 displays a constant degree of relative risk-aversion – it plots as a straight line when both axes use logarithmic scales. For each of the other years, the curve becomes less steep as income increases or, viewed the other way, steeper as income decreases. This will be the case for any lockbox that contains both TIPS and the market portfolio. Why? Because the curve must approach a vertical line (formally, an asymptote) at the level of income provided by the safe asset (in this case, TIPS).

The absolute value of the slope of such a curve is defined as *relative risk aversion.* Thus each of the curves for periods funded by lockboxes with both a market portfolio and a riskless real asset reflects *decreasing relative risk aversion.* For higher levels of income, the recipient changes the chosen amount of income less in states with higher or lower cost (price per chance). Anyone who chooses a combination of a safe asset and the market portfolio has thus made a choice that is consistent with maximizing a utility function with decreasing relative risk aversion.

The result applies more broadly. Consider retirees who are receiving Social Security payments and invest all their other money in the market portfolio. Absent a change in the Social Security rules, their total real income can never fall below their Social Security payments, which provide an *income floor.* In a given year, no matter how large the price per chance may be for a particular state (scenario), their total income will be at least as large as that provided by Social Security. In either of the two previous diagrams, as one moves to higher values of PPC, the curve will become steeper, since it can never cross the vertical line representing Social Security income.

The bottom line is that most retirees choose retirement income strategies consistent with utility functions with decreasing relative risk-aversion. But we know that the market portfolio is consistent with constant relative risk aversion. So, for there to be equilibrium, some people must take positions consistent with increasing relative risk aversion. Rich investors are likely candidates, as are those who have yet to reach retirement age. And also our children and grandchildren, who will have to pay taxes to provide additional funds to finance our Social Security payments after we retire and to cover payments on TIPS and other government securities.

Enough about equilibrium. Let's return to the analysis at hand.

Next, we set *analysis.plotYearlyPVs* to 'y', producing the following figure:

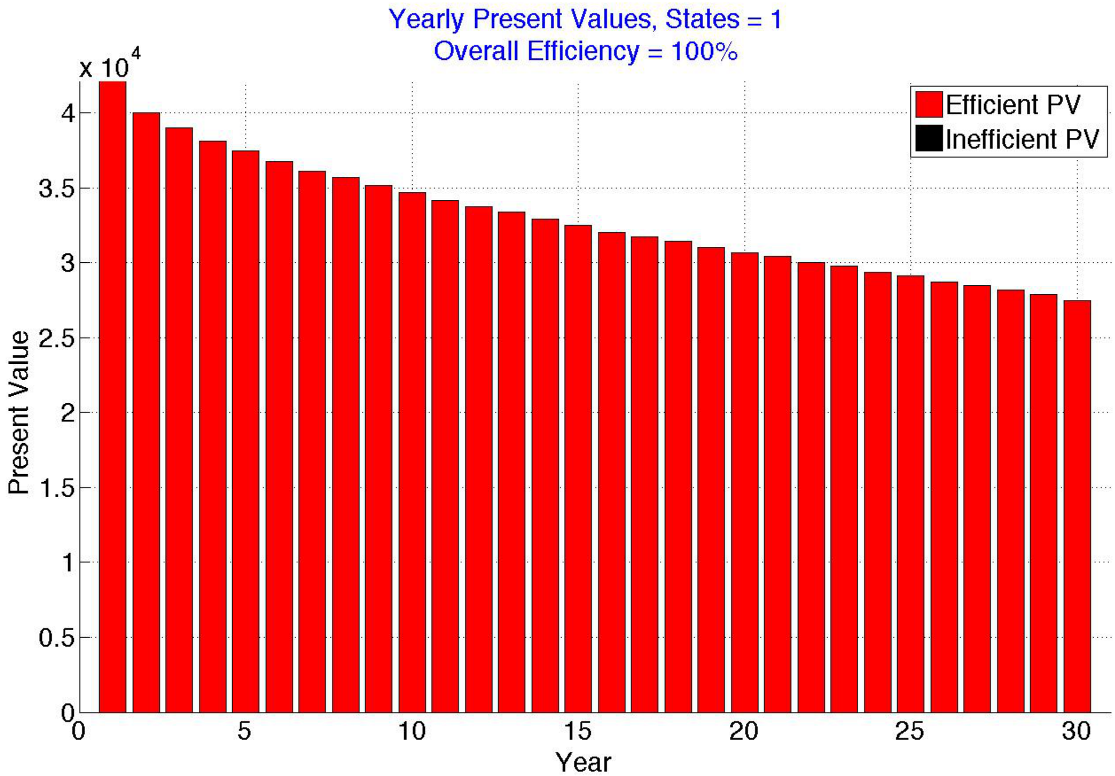

The present values of the incomes produced in each year are very close to the values in our lockboxes (although some could differ slightly due to sampling error). And, as intended, the strategy is completely cost-efficient – there is no way to produce the chosen probability distributions of income at lower cost.

Finally, we set *analysis.plotEfficientIncomes* to 'y' to show again that our strategy is cost-efficient and produces incomes in each year that fall on a linear functions of the cumulative market return. The result of the animation after 29 years have been shown is below. For each year, the actual results fall precisely on a fitted straight line. Moreover, the lines show the payoffs we intended – comforting, if not surprising.

Constant Relative Risk Aversion

As we have seen, a desire to have roughly similar distributions of future income is consistent with different utility functions for each year, each of which exhibits decreasing relative risk aversion as income increases.

On the other hand, we might consider retirees with constant relative risk aversion. One such possibility would involve investing all of the money in each lockbox in the market portfolio and putting the same amount of money in each one. Consider a case with $40,000 invested in each lockbox. The implied marginal utility functions through year 29 are shown below.

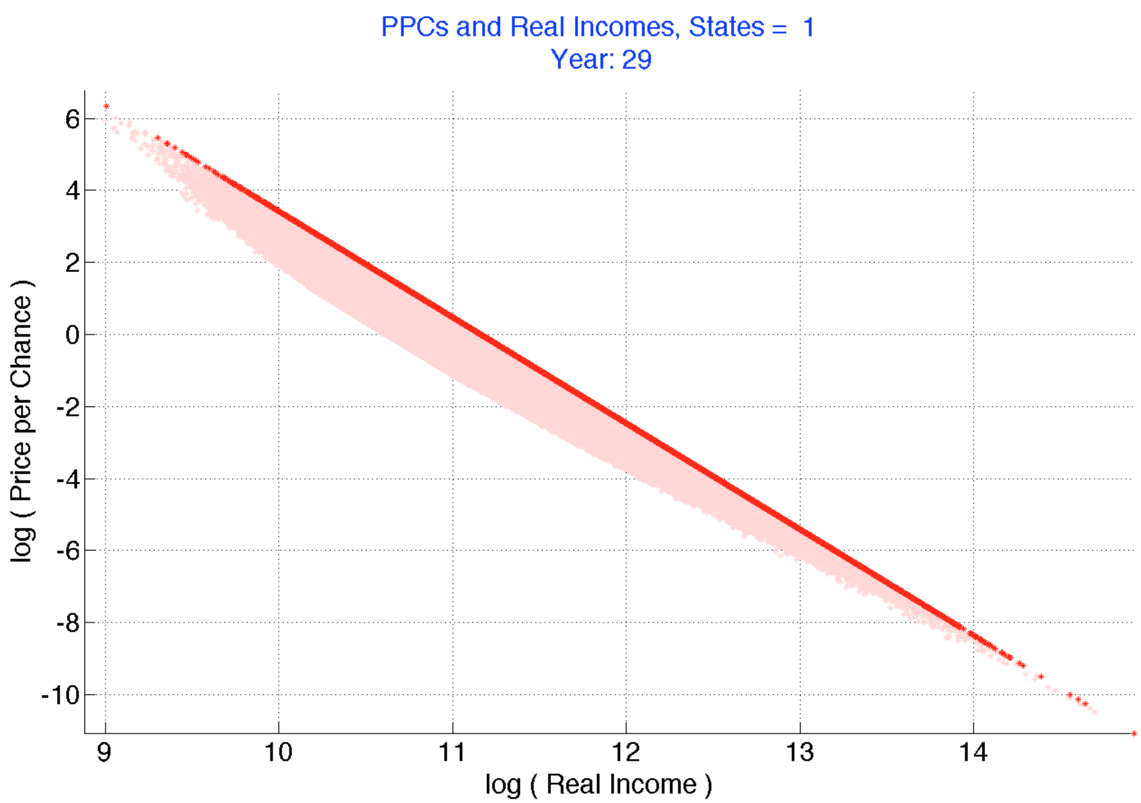

Each function plots as a straight line in a log/log graph and thus exhibits constant relative risk aversion. As one moves to lockboxes for later years, the range of possible incomes increases and the implied utility functions move upward and to the right.

As shown below, the probability distributions of future real income also move to the right and the risk associated with future income is substantially greater, the farther in the future is the year in which a lockbox matures.

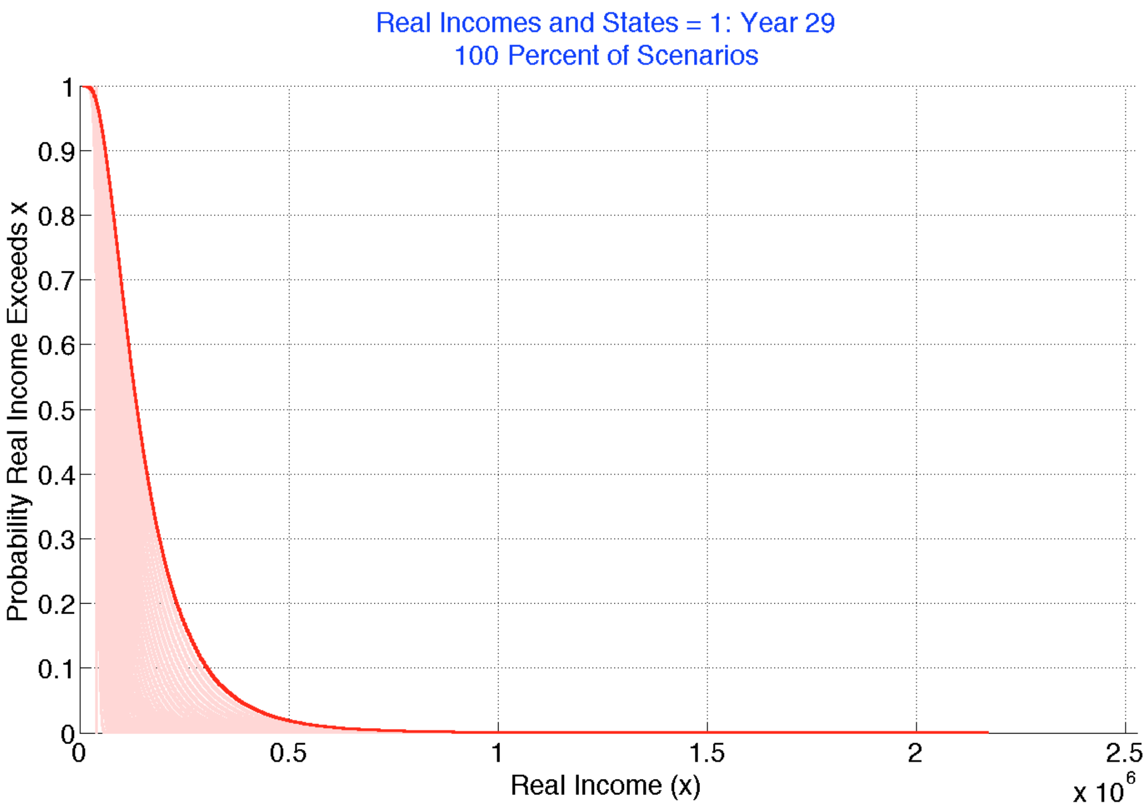

An alternative set of preferences would have the same implied marginal utility function for each future year. This is easily accomplished by putting securities with lower values in lockboxes with later maturities. Analysis of the prior graph of PPCs and Real Incomes can provide the recipe for such a strategy. As we will see later in the chapter, given the parameters we are using for risks and expected returns, the solution is to invest 0.98 times as much in each lockbox as in the one maturing in the prior year.

The graph below shows the implied marginal utility curves for each of the years for a case using this approach with (a) an immediate income (in lockbox 1) of $40,000, (b) $39,200 (= $40,000*0.98) invested in the market portfolio in lockbox 2 (maturing in a year), (c) $38,416 (= $ 40,000 *(0.98^2)) invested in the market portfolio in lockbox 3, and so on.

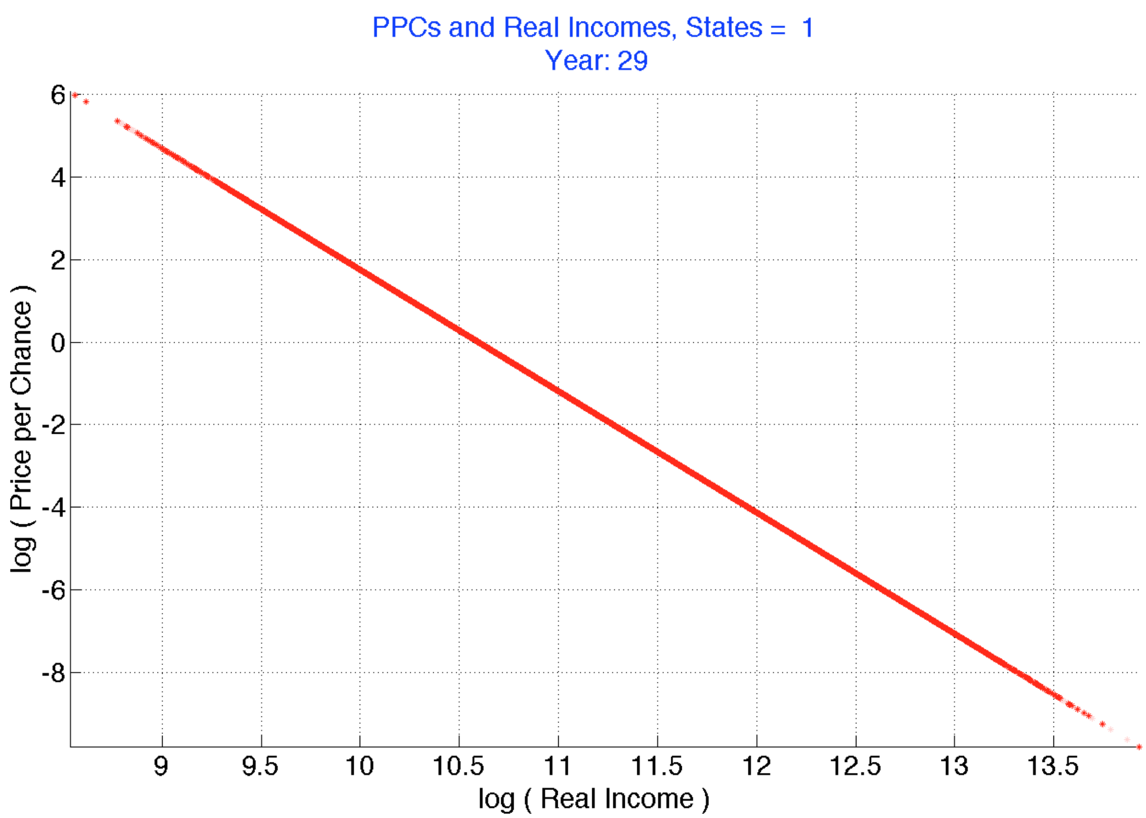

With both axes plotted using logarithmic scales, all the curves are linear with those for later years extending farther towards the axes due to the fact that the longer the holding period, the greater is the range of possible cumulative market returns.

This does indeed reduce the range of possible future incomes, as can be seen in the following graph:

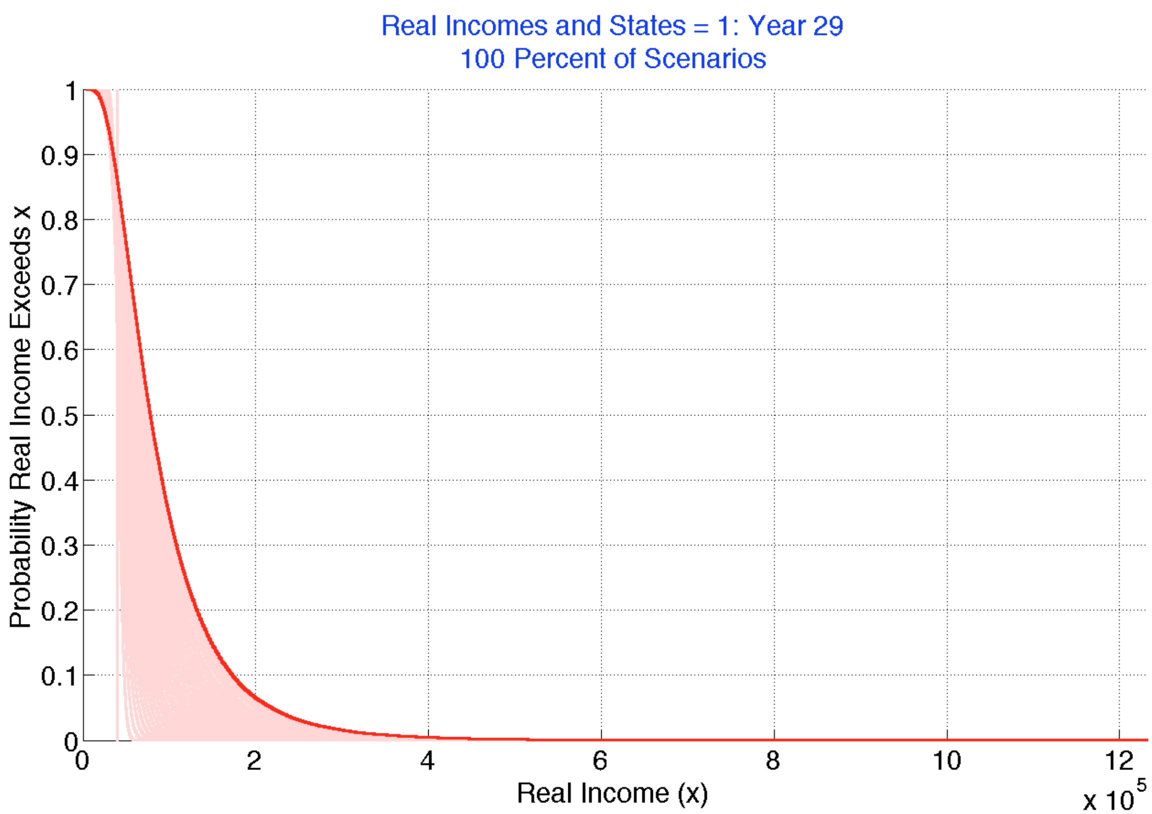

Note the difference in the magnitudes on the horizontal axes of this graph and those in the corresponding one for the prior strategy. In this case the numbers shown on the horizontal axis are to be multiplied by 10^5 while those in the previous version were to be multiplied by 10^6. This is not surprising, since less is invested for every future year.

This may seem strange. Consider an investor with a constant relative risk aversion marginal utility function that spans the entire range in the prior diagram, is the same for every future year, and has the same degree of relative risk aversion as that of the market. One might assume that maximizing utility would result in the same level of risk, somehow defined, for each future year. But most people would likely consider the prospects shown in the graph above to involve greater risk for later years than for future ones. And many would likely prefer some other attainable set of prospects.

But what about a strategy with a constant level of relative risk aversion but one that differs from that of the market as a whole? Perhaps a retiree could select an approach that is optimal for different constant relative risk aversion marginal utility functions for each year, with each function more conservative than the prior one.

Here is an example. Given the parameters we have chosen for expected returns and risks, the relative risk aversion for the market (shown by the slope in the graph with the logarithm of cumulative market return on the x-axis and the logarithm of PPC on the vertical axis) is -2.9428. What about choosing the market portfolio for the first year and a more conservative strategy for a later year? For example consider one for a constant level of relative risk aversion of, say, -2.0.

The figure below shows two such marginal utility functions.

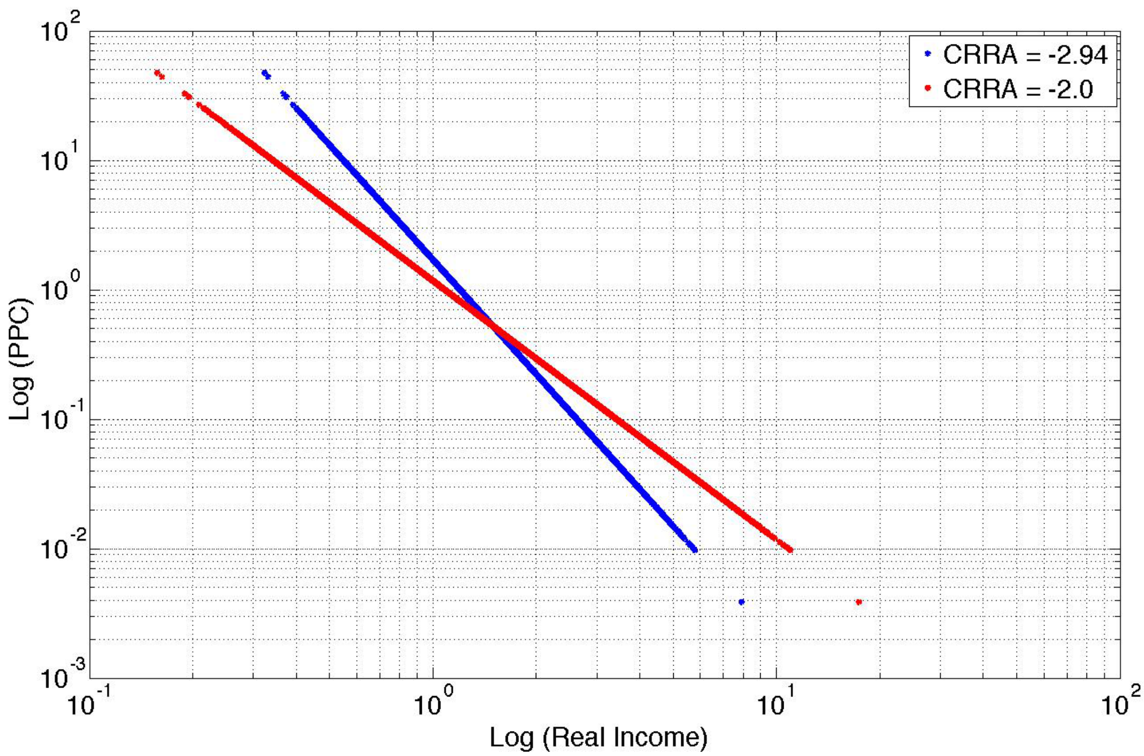

We know that the first function will be consistent with an investment that is wholly invested in the market portfolio. But what about the second? Here is a graph showing the strategy return on the vertical axis and the market return on the horizontal.

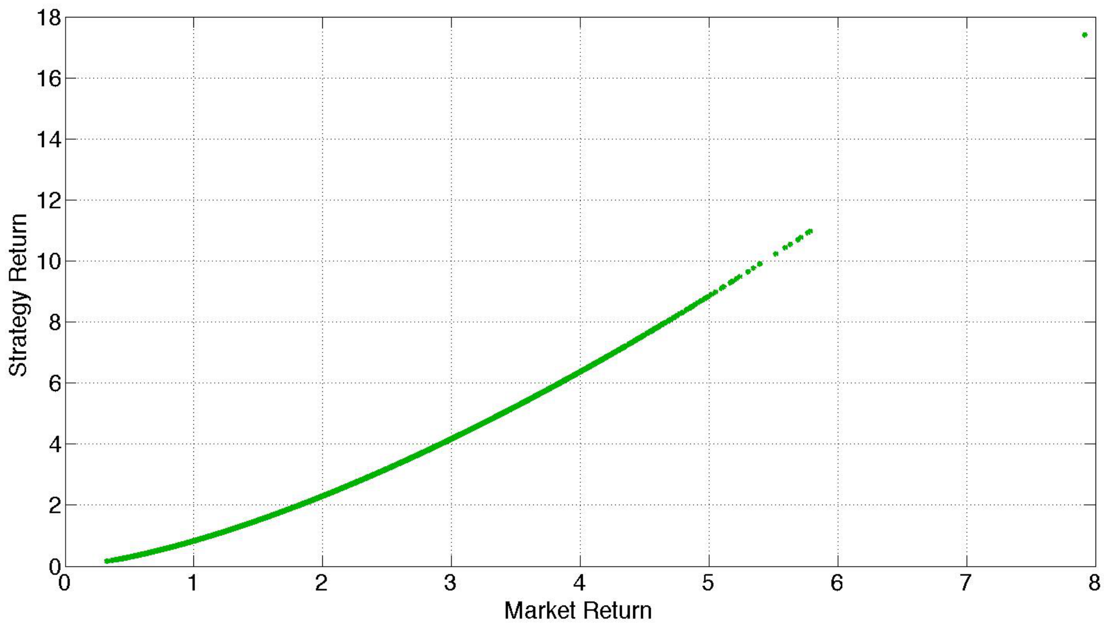

It would be straightforward to produce such range of returns with an m-share, but not with any combination of a risk-free asset and the market portfolio, since this function is non-linear and every possible combination of a risk-free asset and the market portfolio will provide returns that plot as a linear function of market return.

The conclusion is that absent the availability of suitable m-shares, it is impossible to obtain a portfolio optimal for an investor with a constant relative risk aversion marginal utility function with risk aversion that differs from that priced in the market portfolio. That said, it is possible to choose a combination of TIPS and the market portfolio for one year, find the implied marginal utility function, then find combinations for subsequent years that are optimal for the same marginal utility. The next section shows how.

Constant Marginal Utility Lockboxes

Our goal is to determine holdings of TIPS and/or the market portfolio in a series of lockboxes that will produce probability distributions of returns in each year with the same implied marginal utilities of income.

We begin by constructing a data structure *CMULockboxes* for such constant marginal utility lockboxes:

```
function CMULockboxes = CMULockboxes_create( );
 % creates a CMU lockboxes data structure

  % initial lockbox market proportion: 0 to 1.0 inclusive
    CMULockboxes.initialMarketProportion = 1.0;

  % lockbox proportions (computed by CMULockboxes_process)
    CMULockboxes.proportions = [  ];

  % show lockbox proportions (y or n)
    CMULockboxes.showProportions = 'n';

end
```

The key element is the proportion of the first lockbox invested in the market portfolio (with the remainder in TIPS). As with our AMDnLockboxes, the goal is to produce a set of lockboxes with relative proportions of TIPS and the market portfolio. When the data structure is processed, a matrix with the proportions will be placed in the *proportions* element. Given the total amount of money to be invested, the actual values of the securities held in each lockbox can subsequently be computed by multiplying the lockbox proportions by an appropriate constant.

The final element indicates whether or not a bar graph showing the proportions is to be displayed.

The computations are of course handled by a separate function, here *CMULockboxes_process*. The first section handles the computations, the second the plotting (if desired). Here is the first portion:

```
function CMULockboxes = CMULockboxes_process( CMULockboxes, market, client );
  % computes lockbox proportions for an CMULockbox strategy

  % get number of years
    [nscen nyrs] = size( market.cumRmsM );

  % set proportions for year 1
    mktprop = CMULockboxes.initialMarketProportion;
    if mktprop > 1; mktprop = 1; end;
    if mktprop < 0; mktprop = 0; end;
    tipsprop =  1 - mktprop;

  % find ratio of market proportion each year to that for the prior year
    a = market.avec(2);
    b = market.b;
    logk = ( -log(1/a) ) / b;
    k = exp( logk );
  % compute market proportions for all years
    mktprops = mktprop* ( (1/k).^(0:1:nyrs-1) );

  % compute TIPS proportions for all years
    tipsprops = tipsprop * ( (1/market.rf).^(0:1:nyrs-1) );

  % compute lockbox proportions;
     CMULockboxes.proportions = [ tipsprops; mktprops ];
```

The function first finds the number of years for the current case from the market data structure, makes certain that the initial market proportion is between zero and one, then computes the associated proportion in TIPS.

The next section computes the market proportions for all the desired years. As can be seen, the amounts depend on the parameters of the implied marginal utility function for the market portfolio – the *a* value for the first year and the *b* (relative risk aversion) value.

Next the proportions in TIPS for the years are computed. Not surprisingly, these depend solely on the riskless real return (which we assume to be the same for every horizon).

Finally, we create a matrix of the proportions in the same format used in for the *AMDnLockboxes*, with a column for each year, the proportions in TIPS in the top row and the proportions in the market portfolio in the bottom row.

Here is the remainder of the *CMULockboxes_process* function. With only slight changes to accommodate a different name, it is the same as the one shown earlier for the *AMDnLockboxes_process* function.

```
   % plot contents if requested
     if lower( CMULockboxes.showProportions ) == 'y'
       xs =  CMULockboxes.proportions;
       fig = figure;
       x = 1: 1: size(xs,2);
       bar(x, xs', 'stacked'); grid;
       set(gca, 'FontSize', 30);
       ss = client.figurePosition);
       set(gcf, 'Position', ss);
       set(gcf, 'Color', [1 1 1] );
       xlabel( 'Lockbox Maturity Year ','fontsize', 30 );
       ylabel( 'Amount Invested at Inception   ', 'fontsize', 30 );
       legend( 'TIPS ', 'Market ' );
       ax = axis; ax(1) = 0; ax(2) = nyrs+1; ax(3) = 0; ax(4) = 1; axis(ax);
       t = [ 'Lockbox Proportions for Constant Marginal Utility ' ];
       title( t, 'Fontsize', 40, 'Color', 'b' );
       beep; pause;
     end; % if lower(CMULockboxes.showContents) = 'y'

 end % function
```

Lockbox Combinations

Thus far we have provided for two somewhat extreme approaches to the creation of lockbox proportions. The first attempts to generate incomes with approximately similar probability distributions in different years, without taking into account the fact that in more distant years there are greater ranges of cost (price per chance). The second generates incomes that conform with the same marginal utility function, achieving similar responses to differences in cost (price per chance) but generating considerable differences in the probability distributions of income.

It is entirely possible that a retiree (or a couple thereof) might prefer an approach that compromises on the two possible objectives – with asset allocations falling between the two extremes. To accommodate such cases, we provide a data element that can produce a set of lockboxes with contents equal to a weighted average of two or more other sets of lockboxes.

Here is the function for creating such *combinedLockboxes*:

```
function combinedLockboxes = combinedLockboxes_create( );
  % creates a lockbox by combining other lockboxes

    % lockboxes to be combined (data structures)
     combinedLockboxes.componentLockboxes = {     };

    % proportions of lockboxes being combined
    %   one value for each lockbox; values greater than or equal to 0
    %   values will be normalized to sum to 1.0
     combinedLockboxes.componentWeights = [ ];

    % title of combined lockboxes
     combinedLockboxes.title = 'Combined Lockboxes';

    % combined lockboxes proportions produced by combinedLockboxes_process
     combinedLockboxes.proportions = [ ];

    % show combined lockbox contents (y or n)
     combinedLockboxes.showCombinedProportions = 'n';

end
```

The first element should contain a list of lockbox data structures and the second the desired weight assigned to each of them. The next element allows for a title. After processing, the *proportions* element will contain a matrix with the proportions of TIPS and the market portfolio in the lockboxes, using the format in our prior data structures. The last element determines whether or not the process function should provide a bar chart of the resulting proportions.

The function that produces a combined lockboxes data structure, and a bar chart if desired, is straightforward:

```
function combinedLockboxes = combinedLockboxes_process(combinedLockboxes,market, client);
  % combines componentLockboxes in combinedLockboxes to create a new lockbox
    n = length(combinedLockboxes.componentLockboxes);
    wts = combinedLockboxes.componentWeights;
    wts =  max( wts, 0 );
    wts = wts / sum(wts);

    boxprops = combinedLockboxes.componentLockboxes{1}.proportions;
    combprops  = wts(1) * boxprops;
    for i = 2:length( combinedLockboxes.componentLockboxes )
       boxprops = combinedLockboxes.componentLockboxes{i}.proportions;
       combprops  = combprops + ( wts(i) * boxprops );
    end;
    combinedLockboxes.proportions =  combprops;

  % plot contents if requested
     xs =  combinedLockboxes.proportions;
     nyrs = size( xs, 2 );
     if lower( combinedLockboxes.showCombinedProportions ) == 'y'
       fig = figure;
       x = 1: 1: size(xs,2);
       bar( x, xs', 'stacked' ); grid;
       set( gca, 'FontSize', 30 );
       ss = client.figurePosition;
       set(gcf, 'Position', ss );
       set(gcf, 'Color', [1 1 1 ]);
       xlabel( 'Lockbox Maturity Year ', 'fontsize', 30 );
       ylabel( 'Amount Invested at Inception   ', 'fontsize', 30 );
       legend( 'TIPS ', 'Market ' );
       ax = axis; ax(1) = 0; ax(2) = nyrs+1; ax(3) = 0; ax(4) = 1; axis(ax);
       t = [ 'Lockbox Proportions for ' combinedLockboxes.title ];
       title( t, 'Fontsize', 40, 'Color', 'b');
       beep; pause;
     end; %if lower(combinedLockboxes.showContents) = 'y'

end % function
```

The create and process functions could have provided for graduation ratios so that lockboxes for later years would have lower or higher income distributions and corresponding marginal utility functions. Instead, such a feature will be included instead in programs that use lockboxes to provide annuity payments or non-annuitized incomes.

Here are the proportions for a combination with equal proportions of AMD2 and CMU lockboxes, obtained by setting:

combinedLockboxes.componentWeights = [0.5000 0.5000];

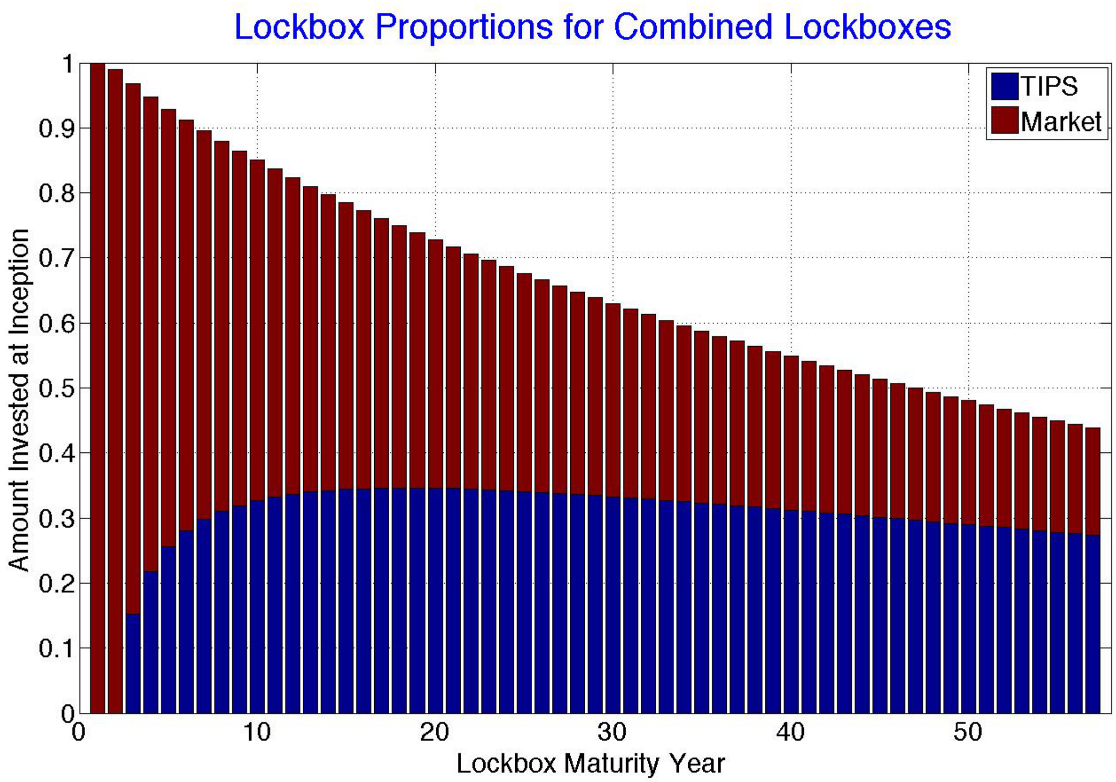

Not surprisingly, the proportion for each year is a 50/50 combination of the proportions for the two other strategies.

Absent the availability of m-shares, retirees can either need to choose retirement income strategies that provide payments with significantly different probability distributions of income at future dates, adopt an approach consistent with marginal utility functions of future income that differ substantially, or select some combination of the two approaches. Moreover, given the fact that most retirees will receive fixed real payments from Social Security or some other sort of defined benefit plan, it will be important take into account all sources of income. The remaining chapters explore some of these implications in detail for the construction of strategies for providing income with and without insuring against longevity risk. In both contexts, lockboxes can play a prominent role.