Retirement Income Analysis
with scenario matrices

William F. Sharpe

18. Proportional Spending

Spending Based on Portfolio Value

The constant spending approach analyzed in Chapter 17 provides real income in year 1 that is a pre-specified proportion of the total value of a portfolio designed to provide retirement income. But in subsequent years, the real amount spent is fixed, bearing no relationship whatever to the value of the portfolio designed to provide income. This chapter deals with a polar opposite – strategies in which spending in each year is a predetermined proportion of the value at the time the portfolio is supplying income. Two key aspects of any such strategy are rules for (1) the asset allocation of the portfolio each year and (2) the proportion of the portfolio value to be spent in each year. We deal first with the latter.

Lockbox Equivalence

Consider a spending rule in which the proportion of portfolio value to be spent in year t is proportion p_t of the portfolio value at the beginning of year t. Let R_{st} be the ratio of the real value for scenario s of the portfolio at the end of year t to the value at the beginning of the year, after the amount to be spent has been withdrawn. The set of all such returns will be a matrix of returns computed (as in the previous chapter) using the combinations of assets for the strategy being analyzed.

Assume that the initial value of the portfolio before the first withdrawal is V. Then the amount withdrawn in each scenario at the beginning of year 1 will be:

$$Vp_1$$

the amount withdrawn in scenario s at the beginning of year 2 will be:

$$V(1-p_1)R_{s1}\,p_2$$

Now, let's rearrange the formula:

$$[V(1-p_1)p_2]\;R_{s1}$$

and again:

$$[(1-p_1)p_2V]\;R_{s1}$$

The bracketed expression can be considered a dollar amount, since it equals the initial value $\boldsymbol{V}$ multiplied by two constants. For example, assume that 5% of the initial portfolio value is to be allocated to spending at the beginning of year 1 (immediately) and 6% of the value of the portfolio value is to be allocated to spending at the beginning of year 2. Thus $p_1 = 0.05$ and $p_2 = 0.06$. Substituting these values, the amount to be spent in scenario s will be:

$$[(1-0.05)*0.06]\,V\,R_{s1}$$

or:

$$[0.057V]R_{s1}$$

Which is equal to the value at the beginning of year two of a portfolio with $0.057V$ dollars invested at the beginning of year 1 times the total return of the portfolio in the first year for the scenario.

Consider next the amount spent at the beginning of year 3 in scenario s. It will be:

$$V(1-p_1)R_{s1}(1-p_2)R_{s2}p_3$$

Rearranging:

$$[(1-p_1)(1-p_2)p_3V]R_{s1}R_{s2}$$

And this will equal the value at the beginning of year 3 of a portfolio with an initial real value equal to the amount given by the bracketed expression times the cumulative returns for scenario s in years 1 and 2.

Now, consider the amounts spent in all the scenarios at the beginning of year 3. Each will equal a value given by the formula above. They will likely differ due to differences across scenarios in the total returns in the two years. But each one will equal the total value of an initial portfolio with the value given by the bracketed expression.

We can continue. The general relationship for spending at the beginning of year T will be:

$$[(1-p_1)(1-p_2)\ldots(1-p_{T-1})p_TV]\ \ R_{s1}R_{s2}\ldots R_{s,T-1}$$

Thus the amount spent in a given year in a scenario will equal the total value of a portfolio with a value equal to the bracketed expression times the cumulative return on the investment strategy up to the year in question for that scenario.

This should seem familiar. The bracketed expression can be considered the initial value of a lockbox for year T containing investments that will provide the returns for each scenario provided by the assets in the lockbox.

If all the scenario returns are generated by a single initial combination of TIPS and/or the market portfolio with no rebalancing thereafter, the implied lockbox will be *market-based* in the sense that the value at maturity will be a non-decreasing function of the cumulative return on the market portfolio. This will guarantee that the payments made in any year will be a non-increasing function of the price per chance. Thus the strategy will be *least-cost efficient* – no approach can provide the same distribution of real income in that year at lower cost. An alternative efficient approach would invest the initial implied lockbox in some sort of m-shares (but as this is written, none are available).

Some advocates of proportional spending recommend a constant asset allocation, usually divided between stocks and bonds; others favor changing bond/stock allocations from year to year, to follow a predetermined schedule such as a glide path. Some advocate decreasing stock proportions over time, others increasing them. In any event, as long as asset proportions at the beginning of each year are specified at the outset, it would be possible to create a set of lockboxes, each with shares of a fund that follows the same dynamic asset allocation strategy. The results would not be cost-efficient, since payments in at least some years would be path-dependent. We of course prefer cost-efficient lockboxes and will follow the convention that the term “lockbox” without a modifier will imply contents that are market-based.

This chapter provides a function for creating a data structure for proportional spending policies and a function for processing such a structure. However a caveat is in order. Since any proportional spending strategy that provides least-cost efficient incomes can be obtained using a more transparent type of lockbox spending approach, for such cases we recommend using doing this explicitly using the functions provided in the chapter 20. That said, there is a substantial literature on proportional spending approaches, well worth reviewing. And the functions provided in this chapter can be used to such approaches with with or without changing asset allocations over time.

IRS Required Minimum Distributions

Most proportional spending policy advocates derive the proportions to be spent in each year from actuarial tables, using either detailed information about retirees' ages, sex, health, etc.. or generic information for typical cases. Common approaches attempt to provide income that has an x% chance of lasting until the beneficiaries are dead, where x is a parameter either chosen by an advisor or incorporated in a standard strategy. A particularly popular choice sets the proportion spent in year t equal to $1/LE_t$, where LE_t is a retiree's "life expectancy" year t. When there are two retirees, detailed actuarial calculations might be used, but this is relatively rare. Moreover, while the term "life expectancy" is often used, it is often taken to mean the median future time at which there is a 50% chance that an individual or couple will be alive. Variants on the approach may use a more conservative value, such as a 75% chance that someone will be alive.

Our functions for income provision using proportional spending approaches will allow the user either (1) to accommodate life expectancy or similar methods by entering a matrix with points on a curve showing the reciprocals of the proportions of portfolio value to be spent in each year or (2) to utilize an approach based on required taxation of portions of certain tax-sheltered savings in the United States (a method favored by some financial advisors).

The U.S. Internal Revenue Service (IRS) allows a number of retirement vehicles to accumulate savings and investment returns thereon without incurring any income taxes until funds are withdrawn. Key examples are Individual Retirement Accounts (IRAs), Rollover Individual Retirement Accounts, 401(k), profit sharing, 403(b), and other defined contribution plans. However, after the beneficial owner of the account reaches age 70 ½, some money must be withdrawn from such funds each year and treated as income subject to Federal income tax. The IRS publishes tables indicating the *required minimum distribution* (RMD) for each year, expressed as a percentage of the year-end account value, for account owners of different types.

There are alternative tables, but most common is the *Uniform Lifetime Table,* which must be used by (1) unmarried account owners, (2) married owners whose spouses are not more than 10 years younger, and (3) married owners whose spouses are not the sole beneficiaries of the tax-deferred accounts. For each age from 70 to 115 (and over), the table provides a *distribution period*, often considered a sort of life expectancy, expressed in years.

The blue dots in the figure below plot the information in the 2016 table. The red curve is a quadratic function fit to the data.

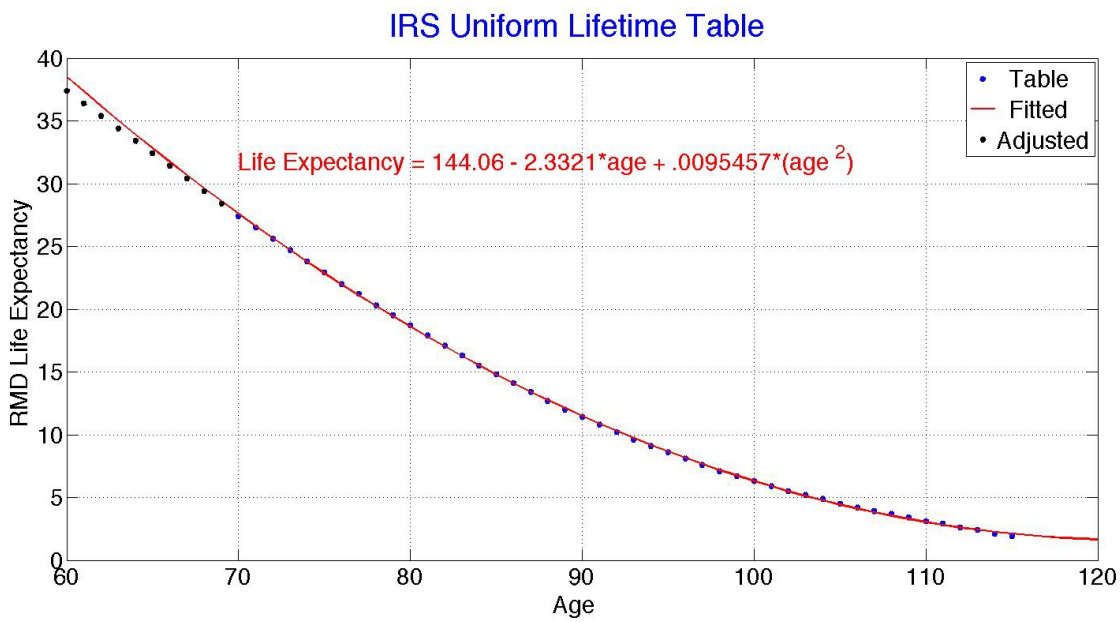

In the figure, the fitted curve is extended to cover ages 60 through 69. However, there is a problem with the result. The life expectancy for age *t-1* cannot more more than 1 year greater than that for age *t*, since the difference would equal 1.0 if there were no chance of dying in the intervening year and less than 1.0 if there were some such chance. Since the fitted curve becomes steeper than 1.0 for ages below 69, we have assumed zero mortality at this and younger ages. The resulting life expectancies are shown by the dots for ages 60 through 69.

The IRS procedure determines the required minimum distribution (RMD) proportion at age 70 or above by simply dividing 1.0 by the life expectancy for that age. The results for the ages covered by the IRS table are shown below, along with those for ages prior to 70 using our calculations.

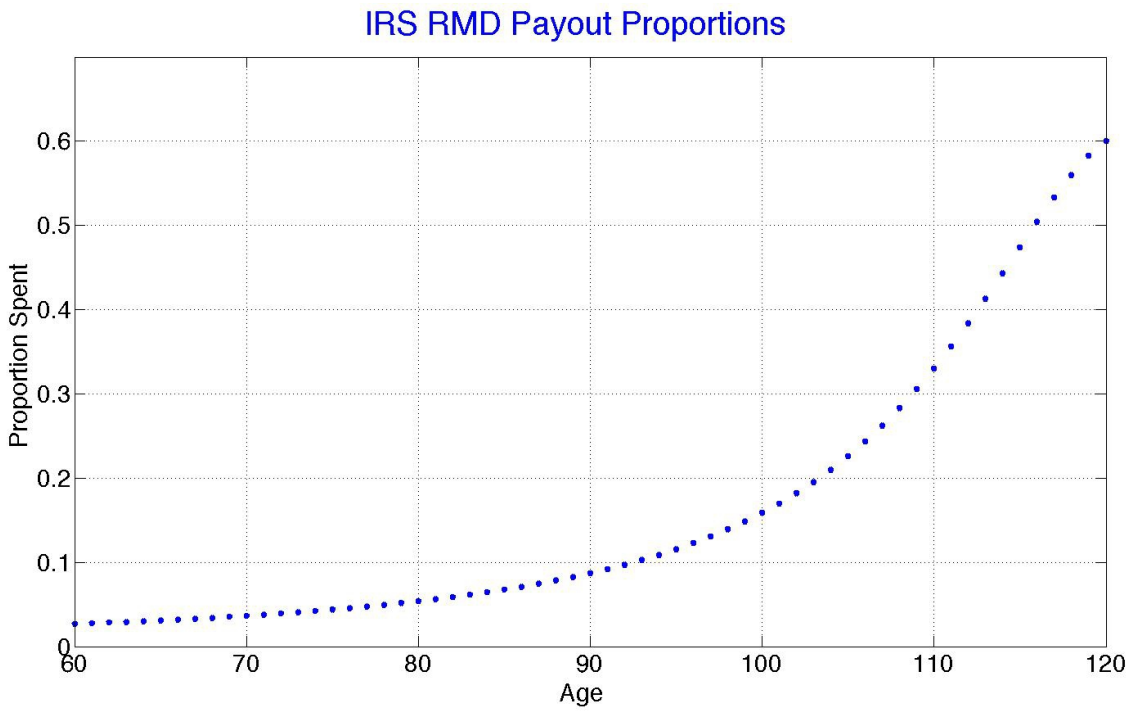

While the IRS life expectancies are undoubtedly based on actuarial tables, they are at best approximations, since they are to be applied whether a beneficiary is single or has a spouse (as long as the latter is not more than ten years younger) and regardless of the latter's age.

Many investors believe that the IRS is recommending that the RMD proportion of covered investment funds be spent each year. But this is not required. To be sure, one must include the associated value in taxable income, but some or all can be retained, put in an investment vehicle in which dividends, interest and capital gains are subject to federal and state income taxes, and spent in later years and/or passed as part of an estate. Nonetheless, a number of retirees choose to spend their required minimum distributions, thus adopting a spending strategy with the proportions shown in the graph above at age 70 and above. A possible extension to the proportions shown for ages 60 through 69 (and, if applicable, similar calculations for earlier ages) would seem a natural corollary.

Our proportional spending functions provide such an approach as an option, and we will use it for the examples in this chapter. The software will, however, accommodate any set of life expectancies (giving corresponding spending proportions) the user may wish to input.

The iPropSpending_create function

The initial part of the function for creating an iPropSpending data structure follows:

```
function iPropSpending = iPropSpending_create( )
 % create a proportional spending income data structure

 % amount invested
    iPropSpending.investedAmount = 100000;

 % use IRS Required Minimum Distributions (RMD) Life Expectancies (y or n)
    iPropSpending.useRMDlifeExpectancies = 'y';
 % if RMD not used, vector of life expectancies and age for first value
    iPropSpending.nonRMDlifeExpectancies = [   ];
    iPropSpending.nonRMDfirstLEAge = 70;

 % current age of portfolio owner
    iPropSpending.portfolioOwnerCurrentAge = 65;

 % show proportions spent (y or n)
    iPropSpending.showProportionsSpent = 'n';

 % show Lockbox equivalent initial investment values
    iPropSpending.showLockboxEquivalentValues = 'n';
```

The first element indicates the value of the account at the present time. The next three provide information needed to compute the proportions of portfolio value to be spent each year. If the RMD life expectancies are to be used, next two data elements are not utilized. Otherwise, a vector of life expectancies should be input, beginning at the indicated age. As with the RMD approach, the life expectancy vector will be expanded, assuming no mortality before the first LE age and constant mortality for ages after the age corresponding to the last element in the life expectancy vector.

The next element shows the current age of the portfolio owner, since it could be either of the two retirees specified in the *client* data structure.

The next element indicates whether or not a graph of the proportions is to be shown. The following element indicates whether or not to show the relative values that would be included in a set of equivalent lockboxes.

The remaining statements of the *iPropSpending* data structure (shown below) are the same as the those used for the *iConstSpending* data structure. The first element allows for portfolio investment to be a constant mix of the market portfolio and bills or to change from year to year. The *iPropSpending.glidePath* element should be a matrix with market proportions in the first row and corresponding years in the second. Proportions for years between points are interpolated, those for years before the first point are the same as that for the first point, and those for years after the last point are the same as that for the last. The next element can be used to show the resulting glide path. The final element provides the *retention ratio*: the proportion of total return that will be available after deducting expenses charged by any involved investment firms and/or advisors.

```
% matrix of points on portfolio market proportion glide path graph
% top row is y: market proportions (between 0.0 and 1.0 inclusive)
%   bottom row is x: years (first must be 1 or greater)
% first proportion applies to years up to and at first year
% last proportion applies to years at and after last year
% proportions between two years are interpolated linearly
   iPropSpending.glidePath = [ 1.0 ;  1 ];

% show portfolio glide path (y or n)
    iPropSpending.showGlidePath = 'n';

% retention ratio for portfolio investment returns
%   = 1 - expense ratio
%   e.g. expense ratio = 0.10% per year, retentionRatio = 0.999
    iPropSpending.retentionRatio = 0.999;

end
```

The next section, on the *iPropSpending_process* function, provides the details of the statements that use these elements to produce matrices of incomes and fees, then add this information to the corresponding client matrices. Those not fascinated by MATLAB code may wish to read only the descriptive text or trust that the function does its tasks appropriately.

The iPropSpending_process function

The initial statements in this function use elements in an *iPropSpending* data structure to create a complete matrix for the investment glide path; they are the same as those in *iConstSpending_process* function:

```
function client = iPropSpending_process( iPropSpending, client, market );

 % get matrix dimensions
   [ nscen nyrs ] = size( market.rmsM );
 % get glidepath
   path =  iPropSpending.glidePath;
 % get points from glidepath
  ys =  path( 1, : );
  xs =  path( 2, : );
 % insure no years prior to 1
  xs = max( xs, 1 );
 % insure no market proportions greater than 1 or less than 0
  ys =  min( ys, 1 );
  ys =  max( ys, 0 );
 % sort points in increasing order of x values
  [ xs ii ] = sort( xs );
  ys = ys( ii );
 % add values for year 1 and/or last year if needed
  if xs(1) > 1;  xs = [ 1 xs ];  ys = [ ys(1) ys ]; end;
  if xs( length(xs) )< nyrs
     xs = [ xs nyrs ]; ys = [ ys ys(length(ys)) ];
  end;

 % create vectors for all years
  pathxs = [ ]; pathys = [ ];
  for i = 1: length( xs ) - 1
    xlft = xs(i);  xrt =  xs(i+1);
    ylft = ys(i);  yrt =  ys(i+1);
    pathxs = [ pathxs xlft ];
    pathys = [ pathys ylft ];
    if xlft ~= xrt
      slope = ( yrt – ylft ) / ( xrt – xlft );
      for x = xlft+1: xrt-1
        pathxs = [ pathxs x ];
        yy = ylft + slope * ( x – xlft );
        pathys = [ pathys yy ];
      end; %  for x = xlft+1:xrt-1
    end; % if xlft ~= xrt
  end; % for i = 1:length(xs)-1
  pathxs =[ pathxs xs(length(xs)) ];
  pathys =[ pathys ys(length(ys)) ];
```

The next statements, which show the glide path if desired, are also the same as those in the *iConstSpending_process* function:

```
% show glide path if desired
 if lower( iPropSpending.showGlidePath ) == 'y'
   fig = figure;
   set( gca, 'FontSize', 30 );
   ss =client.figurePosition;
   set( gcf, 'Position', ss );
   set( gcf, 'Color', [1 1 1] );
   xlabel( 'Year ', 'fontsize', 30 );
   ylabel( 'Proportion in Market Portfolio   ', 'fontsize', 30 );
   plot( path(2,:), path(1,:), '*b', 'Linewidth', 4 );
   hold on;
   plot( xs, ys, '-r', 'Linewidth', 2 );
   legend( 'Input ', 'All ' );
   ax = axis;  ax(1) = 0;  ax(2) = nyrs+1;  ax(3) = 0;  ax(4) = 1;  axis(ax);
   t = [ 'Glide Path: Market Proportions by Year  ' ];
   title( t, 'Fontsize', 40, 'Color', 'b' );
   plot( xs, ys, '-r', 'Linewidth', 2 );
   grid;
   hold off;
   xlabel( 'Year ', 'fontsize', 30 );
   ylabel( 'Proportion in Market Portfolio   ', 'fontsize', 30 );
   beep; pause;
   % create blank screen
    figblank = figure; set( gcf, 'Position', ss );
    set( gcf, 'Color', [1 1 1] );
 end; %  if lower(iPropSpending.showGlidePath) == 'y'
```

Next, as before, we create a matrix with gross returns for the investment strategy in each year:

```
% create matrix of gross returns for investment strategy
 retsM = zeros( nscen, nyrs );
 for yr = 1: nyrs - 1
   rets = pathys( yr ) * market.rmsM( :, yr );
   rets = rets + ( 1-pathys(yr) ) * market.rfsM( :, yr );
   retsM(:,yr) = rets;
 end;
```

and get the retention ratio:

```
% get retention ratio
  rr = iPropSpending.retentionRatio;
```

Finally we turn to the features that distinguish this spending approach from others. If the RMD life expectancies are to be used, we initialize vectors with the RMD values and the first age for which they apply. Otherwise we obtain a vector of life expectancies and the applicable first age from the corresponding elements of the *iPropSpending* data structure:

```
% get life expectancies
 if lower( iPropSpending.useRMDlifeExpectancies ) == 'y'
  LEs = [ 27.4 26.5 25.6 24.7 23.8 22.9 22.0 21.2 20.3 19.5 18.7 17.9 17.1 ...
     16.3 15.5 14.8 14.1 13.4 12.7 12.0 11.4 10.8 10.2 9.6 9.1 8.6    ...
     8.1 7.6 7.1 6.7 6.3 5.9 5.5 5.2 4.9 4.5 4.2 3.9 3.7 3.4 3.1  ...
     2.9 2.6 2.4 2.1 1.9 ];
  firstLEAge = 70;
 else
   % if RMD not used, vector of life expectancies and age for first value
   LEs  = iPropSpending.nonRMDlifeExpectancies;
   firstAge = iPropSpending.nonRMDfirstLEAge;
 end; %  if lower(iPropSpending.useRMDlifeExpectancies) == 'y'
```

Next we expand the life expectancy vector to cover all possible ages that might be needed, assuming that there is no mortality before the first given age and that life expectancies are constant from the last given age onward:

```
% expand LE vector
% assume no mortality before first age
  firstLE = Les( 1) ;
  initLEs = firstLE + ( firstLEAge-1: -1: 1 );
% assume life expectancy constant after last age
  LEs = [ initLEs  Les ];
  LEs = [ LEs  Les(length(LEs))*ones( 1, 120 ) ];
% set life expectancies for years based on owners current age
 currAge = iPropSpending.portfolioOwnerCurrentAge;
 LEs = LEs( currAge: length(LEs) );
 LEs = LEs( 1: nyrs );
```

The resulting vector of life expectancies is then used to create a vector of the proportions of portfolio value to be spent in each year, then guaranteeing that all values lie between 0 and 1 inclusive:

```
% find spending proportions and insure they are between 0 and 1 inclusive
  spendProps = 1 ./ LEs;
  spendProps = max( spendProps, 0 );
  spendProps = min( spendProps, 1 );
```

The next set of statements provides an optional bar graph showing the proportions to be spent:

```
% if desired, show proportions spent
  if lower( iPropSpending.showProportionsSpent ) == 'y'
    fig2 = figure;
    set( gca, 'FontSize', 30 );
    ss = client.figurePosition;
    set( gcf, 'Position', ss );
    set( gcf, 'Color', [1 1 1] );
    xs = 1: 1: nyrs;
    ys = spendProps;
    plot( xs, ys, '-*r', 'Linewidth', 2 );
    t = [ 'Proportions of Portfolio Spent   ' ];
    title( t, 'Fontsize', 40, 'Color', 'b' );
    xlabel( 'Year ', 'fontsize', 30 );
    ylabel( 'Proportion of Portfolio Value Spent   ', 'fontsize', 30 );
    grid;
    beep;  pause;
  end; %if lower(iPropSpending.showProportionsSpent) == 'y'
```

At this point everything needed to compute incomes and fees is available. However, in deference to users curious to see the equivalent values that could be placed in lockboxes, the function uses the formulas that we derived earlier in the chapter to produce a bar chart, if desired:

```
% if desired, show Lockbox Equivalent Values
  if lower(iPropSpending.showLockboxEquivalentValues) == 'y'
    % find lockbox equivalent values
      facs = 1 - spendProps;
      facs = [ 1 facs ];
      facs = facs( 1: length(facs) – 1 );
      lbVals = cumprod(facs) .* spendProps;
      lbVals = lbVals* iPropSpending.investedAmount;
      fig3 = figure;
      set( gca, 'FontSize', 30 );
      ss = client.figurePosition;
      set( gcf, 'Position', ss );
      set( gcf, 'Color', [1 1 1] );
      bar( lbVals, 'r', 'Linewidth', 2 );
    %   ax = axis; ax(4) = 1; axis(ax);
      t = [ 'Lockbox Equivalent Initial Values  ' ];
      title( t, 'Fontsize', 40, 'Color', 'b' );
      xlabel( 'Year ', 'fontsize', 30 );
      ylabel( 'Lockbox Equivalent Initial Value   ', 'fontsize', 30 );
      grid;
      beep; pause;
    % create blank screen
      figblank = figure; set( gcf, 'Position', ss );
      set( gcf, 'Color', [1 1 1] );
  end; %if lower(iPropSpending.showLockboxEquivalentValues) == 'y'
```

Once again, it is useful to restate the caveat that such lockboxes would contain a fund that adjusts holdings from year to year if needed to conform with any changing proportions in the *iPropSpending.glidePath* matrix. Only if the matrix calls for investment throughout the period to be wholly invested in the market portfolio or wholly invested in TIPS would the lockboxes provide incomes that are a function solely of the return on the market portfolio and hence a function solely of price per chance and thus cost efficient (providing the distribution of income in each year at the lowest possible cost).

Finally (!) we come to the section of the function that creates matrices of incomes and fees to be added, respectively, to the client incomes and fees matrices. We use the spending proportions directly (rather than the equivalent lockboxes), deducting incomes and fees to be paid each year from portfolio values while one or both clients is/are alive (personal states 1, 2 or 3), and paying the remaining portfolio value to the estate in the year after the last client dies (personal state 4). Since our matrices extend to the last possible year in which an estate is paid, any and all portfolio values will be paid out during the years covered in the matrices.

```
% create vector of initial portfolio values
   portvals = ones( nscen, 1 ) * iPropSpending.investedAmount;

% initialize desired spending matrix
   desiredSpendingM = zeros( nscen, nyrs );
 % initialize incomes and fees matrices
   incsM = zeros( nscen, nyrs );
   feesM = zeros( nscen, nyrs );
 % compute incomes paid at the beginning of year 1
   incsM( :, 1 ) =  portvals * spendProps( 1 );
 % compute portfolio values after income payments
   portvals =  portvals – incsM( :, 1 );

 % compute incomes and fees paid at beginning of each subsequent year
  for yr = 2: nyrs
    % compute portfolio values before deductions
      portvals = portvals .* retsM( :, yr-1 );
    % compute and deduct fees paid at beginning of year
      feesV = (1 – rr) * portvals;
      feesM( :, yr ) = feesM( : ,yr ) + feesV;
      portvals = portvals - feesV;
    % compute incomes paid out at beginning of year in states 1,2 or 3
      v = ( client.pStatesM(:,yr) > 0 ) & ( client.pStatesM(:,yr) < 4 );
      incsM(:, yr) =  v .* ( portvals * spendProps(yr) );
    % pay entire value if state 4
      v = ( client.pStatesM(:,yr) == 4) ;
      incsM(:,yr) = incsM(:,yr) + v.*portvals;
    % deduct incomes paid from portfolio values
      portvals = portvals - incsM(:,yr);
  end;
```

Our work done, we add the results to the client incomes and fees matrices, and end the function:

```
 % add incomes and fees to client matrices
   client.incomesM = client.incomesM + incsM;
   client.feesM = client.feesM + feesM;

end
```

Proportional Spending from a Market Portfolio

Our first case assumes that the portfolio with $1,000,000 invested entirely in the market portfolio throughout the years until all the proceeds have been distributed. While the animated graph with each year's distribution provides more detail, for convenience we use income maps for the examples in this chapter. Here is the one for this case:

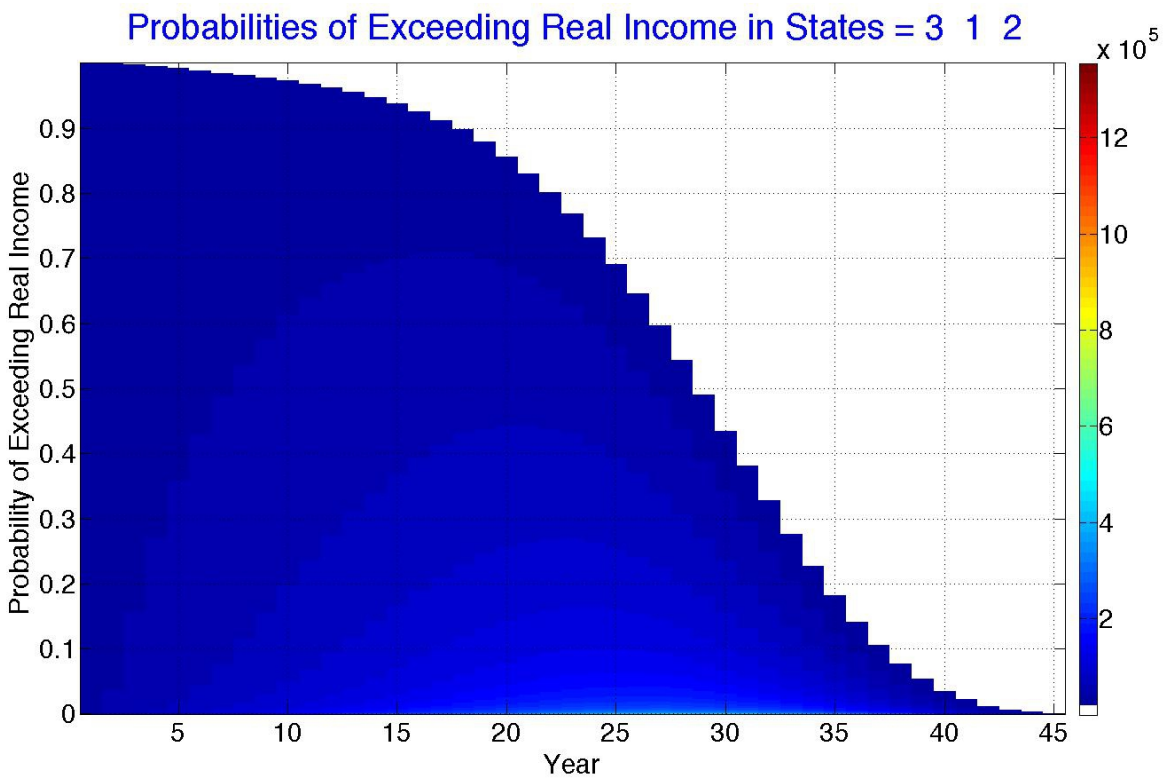

Note that the ranges of income for the early years are relatively low (dark blue). Then in later years there are chances for considerably higher income (light blue) as well.

This can be seen, although less clearly, in the income distribution graph. A frozen version is shown below. The vertical line shows the distribution of roughly $30,000 in year 1. The curves for subsequent years plot farther and farther to the left for high probabilities and to the right for low probabilities until year 28. Those for years after year 29 plot farther and farther to the left. The darker curve, for year 40, reflects the considerably lower possible incomes for later years.

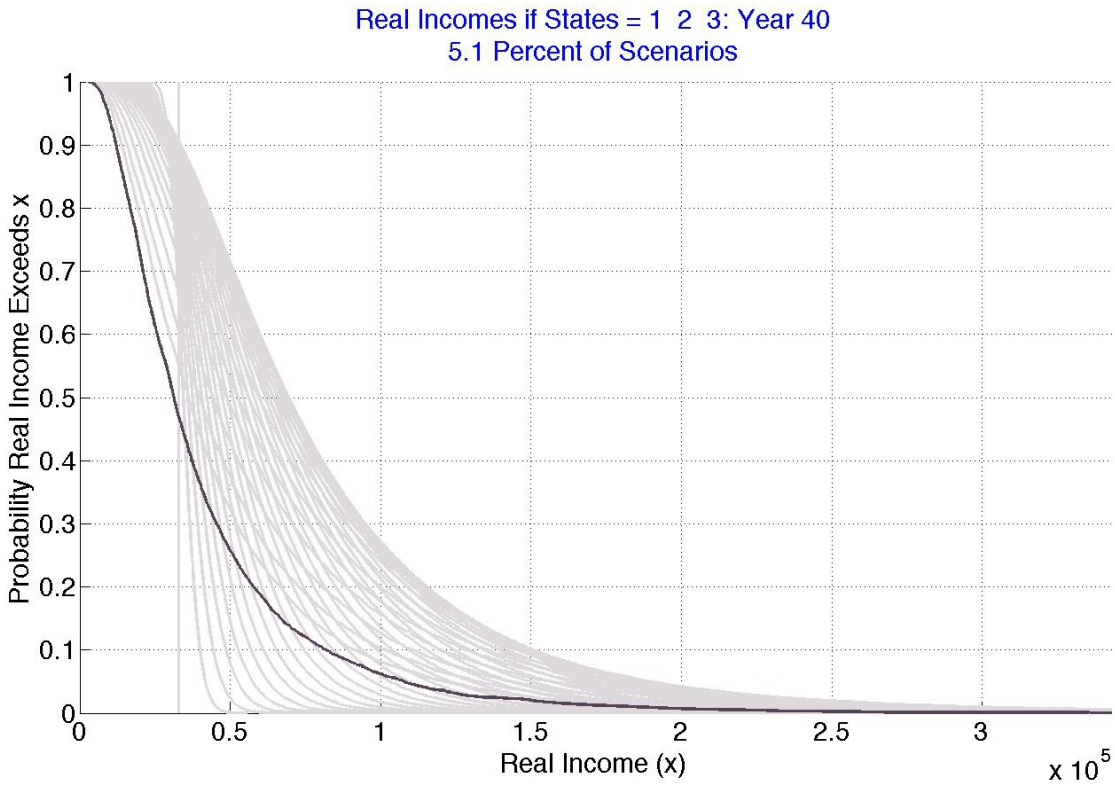

The graph showing PPCs and real incomes tells a story with some of the same features. For convenience we again take a snapshot after an arbitrary year has been plotted. Since the investments are cost-efficient, for each year the points plot on a monotonic downward-sloping curve. Initially, each curve plots slightly to the right of that for the prior year. And, of course, each curve covers a wider range of PPC values. But for given values of PPC, the differences in the points on the curves are relatively small. However, after year 28, the curves begin to plot considerably farther to the left, as indicated by the darker curve for year 40.

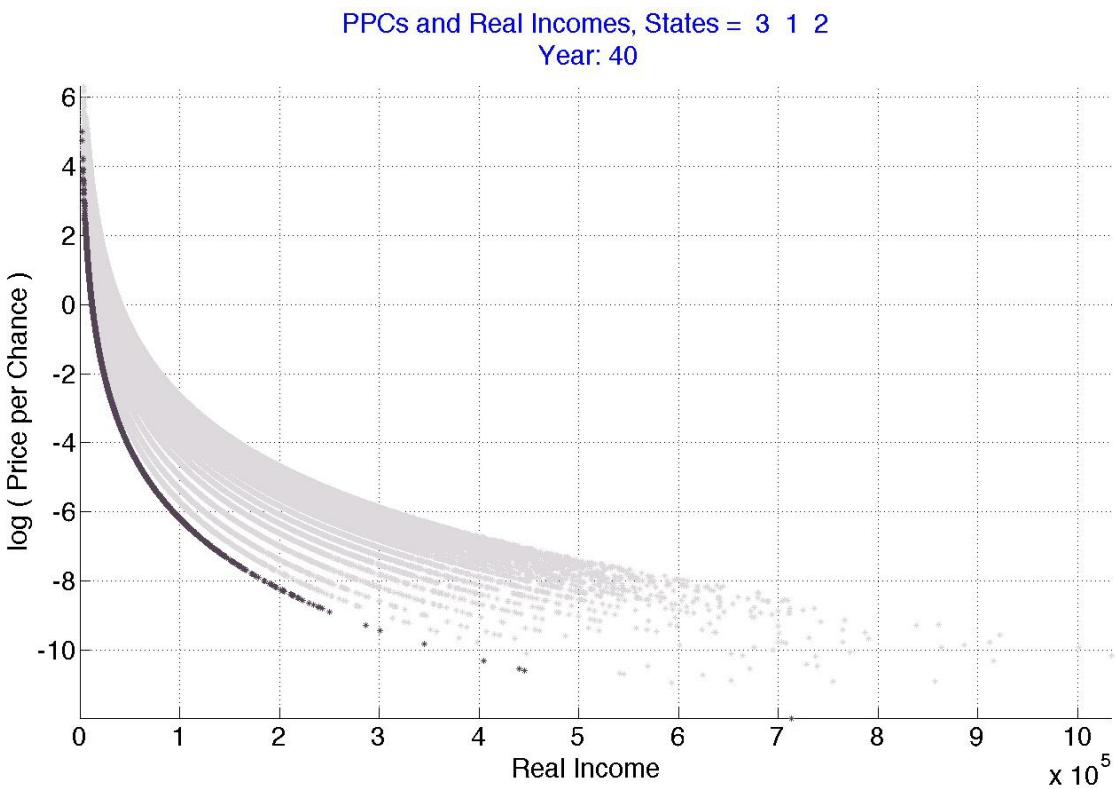

Recall our argument that each such curve can be interpreted as reflecting the implied marginal utility of income for a year (more precisely: plotting on the y-axis the implied marginal utility of income in that year times a constant). However, this would only be true if Bob and/or Sue would be alive for at least 40 years in all the possible scenarios. But such is not the case. It may be reasonable to plan for smaller incomes in distant years if the chances of being alive to enjoy such incomes are small. We will have much more to say about this in chapter 20.

Finally, there is the matter of the distribution of Bob and Sue's savings. The pie chart showing recipient present values tell the tale:

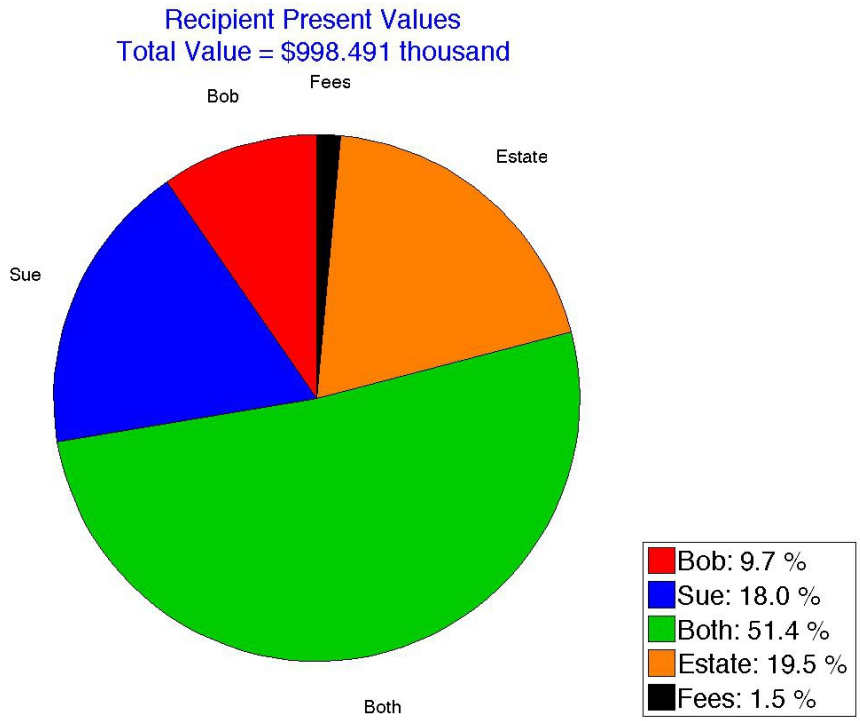

As usual, the sum is close to the amount invested (here, $1,000,000) but differs slightly due to sampling error. The first three wedges tell the usual story that the value of incomes likely to be received when they are both alive will be greatest, followed by the value of the incomes likely to be received when the younger and female Sue is alone, with the still smaller value of the incomes that the older and male Bob might receive if Sue dies first.

The impact of investment fees is mercifully relatively modest, since we assumed investments with low expense ratios: 1/10 of 1% (10 basis points).

The possibly shocking result is the present value of the possible amounts that could go to Bob and Sue's estate: almost 20% of their investment. Even though they chose spending proportions that would provide relatively smaller incomes if they lived very long lives, there were many possible scenarios in which they would leave substantial estates.

This provides a very graphic reminder that without annuities, one has to accept a tradeoff. The smaller the chance of “running out of money”, the greater the likelihood that a substantial amount of money will be left unspent, then provided to an estate.

Proportional Spending from a TIPS Portfolio

Now let's consider another possible cost-efficient approach for a proportional spending strategy – investing the entire portfolio in TIPS for each future year.

The figure below shows one possible scenario, in which Sue outlives Bob and has a long life thereafter.

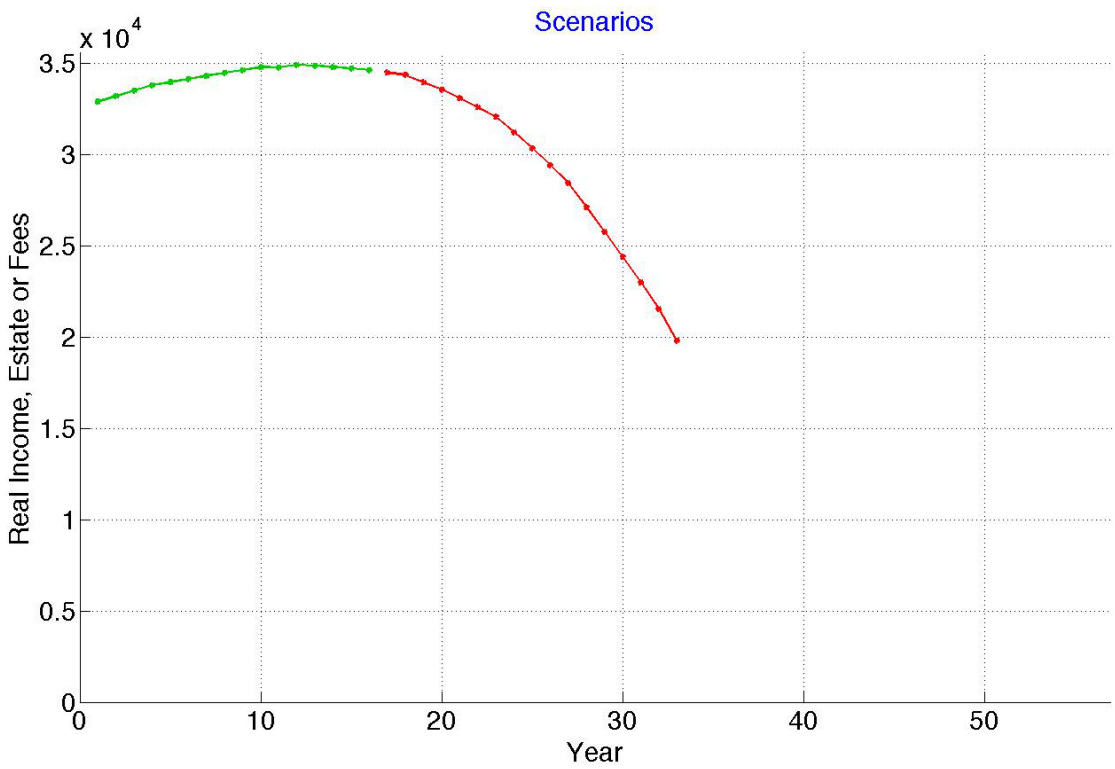

As can be seen, real income increases slightly over the first 20 years, from somewhat over $32,000 to almost $35,000 per year. After that it begins to fall, slowly at first, then more rapidly, reaching slightly less than $20,000 in year 33. In this scenario, Sue then departs leaving the remaining portfolio value to the estate.

According to our market assumptions, there is no uncertainty about the future values of a TIPS portfolio. Hence every scenario will plot along some or all of this curve and its extension, depending on Bob and Sue's lifespans.

The results for all 100,000 scenarios are summarized more colorfully in the income map:

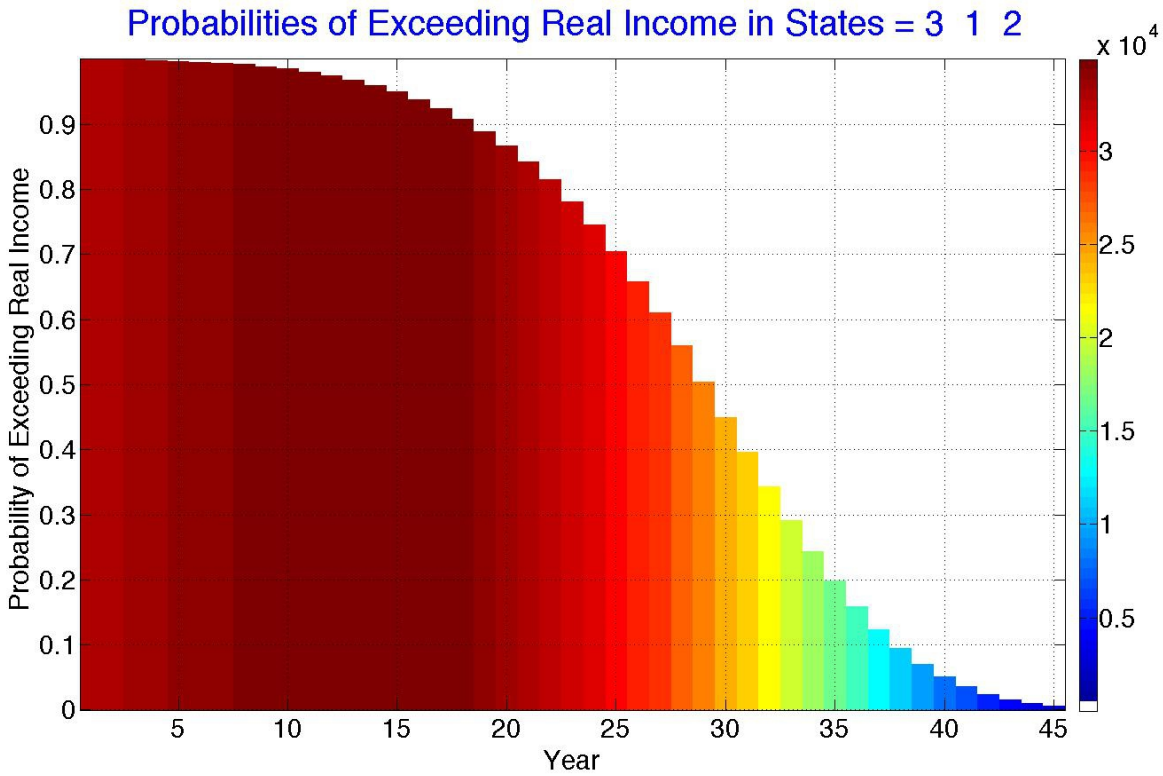

Note that the bar for each year is all one color, indicating no differences in incomes for that year across scenarios. But the colors show the pattern of increasing, then decreasing incomes as time goes on (until the estate is paid) in every scenario.

Now to the present values of the various claims on possible future income from the portfolio. Here are the results:

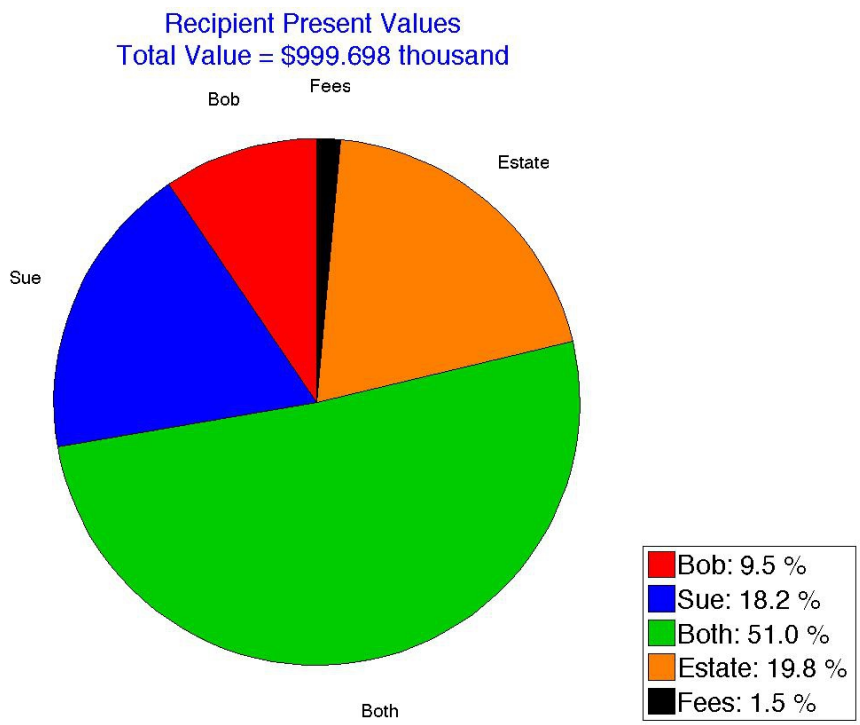

Perhaps surprisingly, the relative percentages for the five claimants are almost the same as those in the prior example, although the ranges of real incomes in each year are very different indeed! Why is this so? The Lockbox Equivalence principle provides the answer. The contents of the lockbox providing income in a year *t* will be allocated according to the personal states in that year. And the initial value of that lockbox is the same, regardless of the manner in which the money in the lockbox is invested. Some recipient (Bob, Sue, Bob&Sue or the estate) will receive the contents of the lockbox in each scenario. Thus the present value should equal the initial value, no matter what the investment strategy. Moreover, the particular recipient for each scenario is uncorrelated with the cumulative returns on the lockbox investments. To be sure, the chance of early receipt by the estate in the event that both Bob and Sue have been gone for more than a year slightly complicates the argument. But for a proportional spending strategy or its lockbox equivalent, the distribution of the initial present value among the possible claimants will be relatively invariant to the investment strategy chosen.

Proportional Spending from a Combination of TIPS and the Market Portfolio

A number of advocates for proportional spending suggest investment of 60% of the retirement fund in a diversified stock portfolio and 40% in some sort of bond fund, with funds reallocated periodically (typically annually) to restore the proportions to 60/40. As we have indicated, this amounts to a strategy of periodically selling relative winners and buying relative losers. And not every investor can do this (with whom would they trade?). Our market portfolio includes stocks and bonds and might thus have risk relatively similar to that of such a 60/40 mix. But our portfolio holds all risky securities in market proportions at all times and could be utilized by all investors. More formally – it can be *macroconsistent*.

In our simplified world, only policies that provide income at any future time t which are a non-decreasing function of the cumulative return on the market portfolio are cost-efficient, in that the associated probability distribution of income cannot be obtained at a lower cost. The two prior examples (100% investment in the market portfolio and 100% investment in TIPS) are efficient in this sense. The strategies in this section and the next two may not be perfectly efficient; thus we will use an analysis computation to measure the degrees of their inefficiencies.

We start with a case in which the retirement fund is invested in a combination with 50% invested in the market portfolio and 50% in TIPS initially, with the fund rebalanced each year to return to a 50/50 value combination.

It is a simple matter to analyze such a strategy. One simply sets:

iPropSpending.glidePath = [0.50 ; 1];

so that the proportion in the market portfolio will be 0.50 at the beginning of each year.

Here are the annual distributions of income for the scenarios in which Bob and/or Sue are alive:

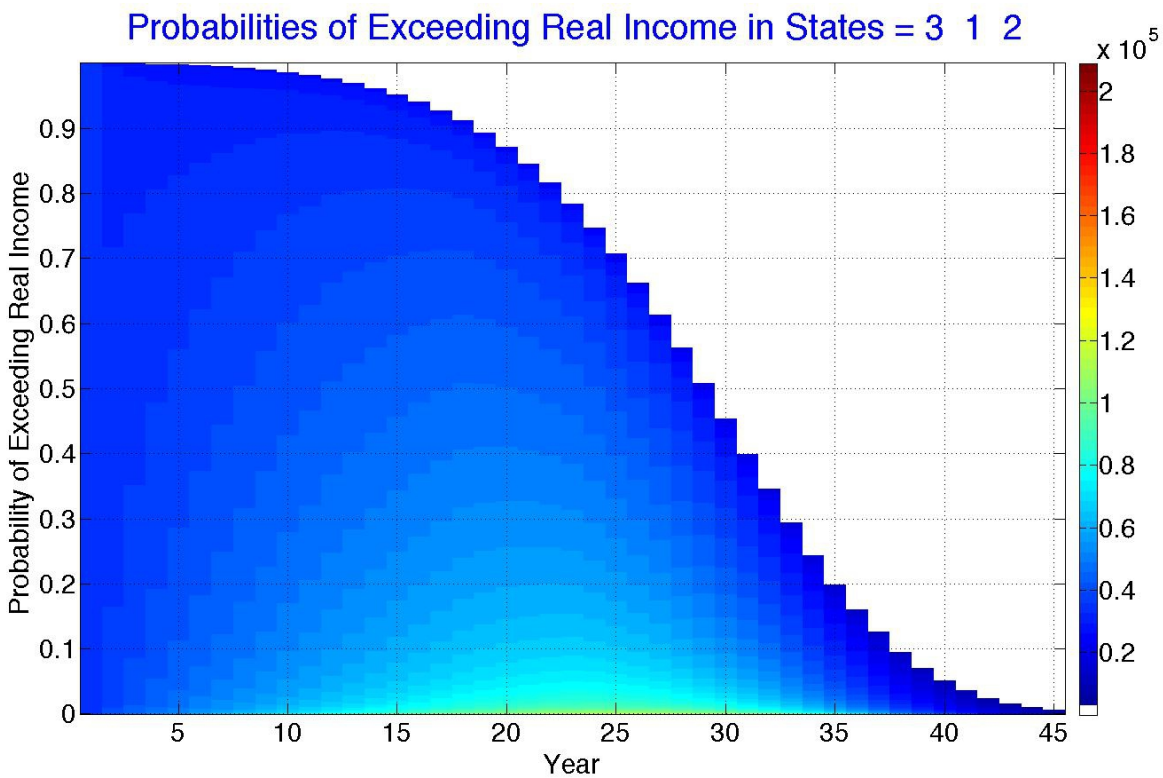

As can be seen, the range of incomes expands from year to year up to year 20 or so. It then begins to contract as the impact of the more miserly implicit initial lockbox values for later years is felt.

The figure below shows that, as anticipated, the distribution of present values is very close to that in the previous two cases.

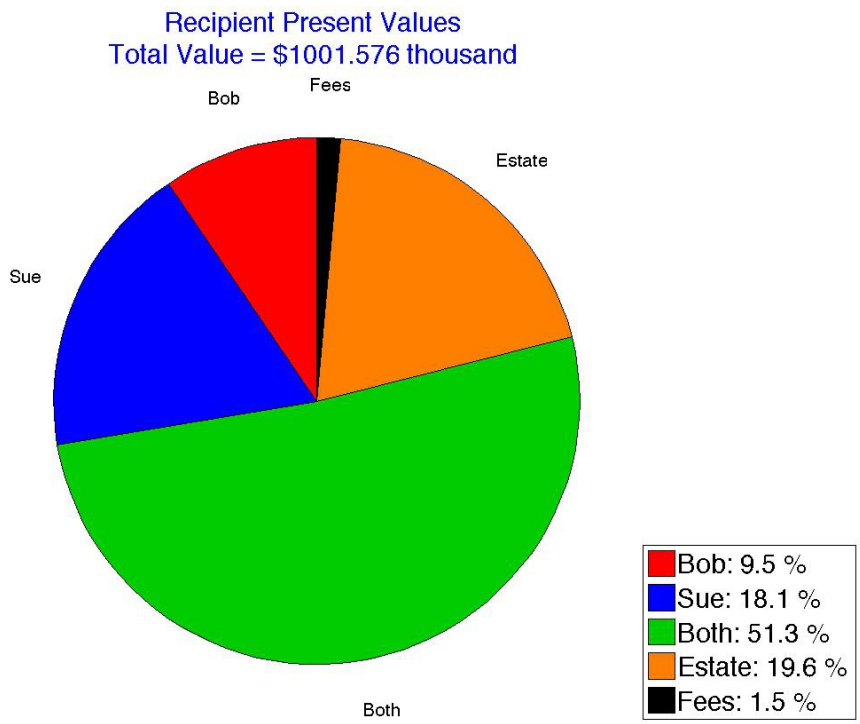

The total value differs slightly from that in the previous cases and from the actual amount, due (as we now know) to sampling error.

Turning to implied marginal utility, we find that the plots for PPCs and Real incomes vary from year to year in a manner similar to that in the prior examples. However, where previously each year's plot was a neat curve, here there is some scatter around the central tendency for each one, as shown here for year 20 (the darker dots). This, of course, reflects the inefficiency of a rebalancing strategy.

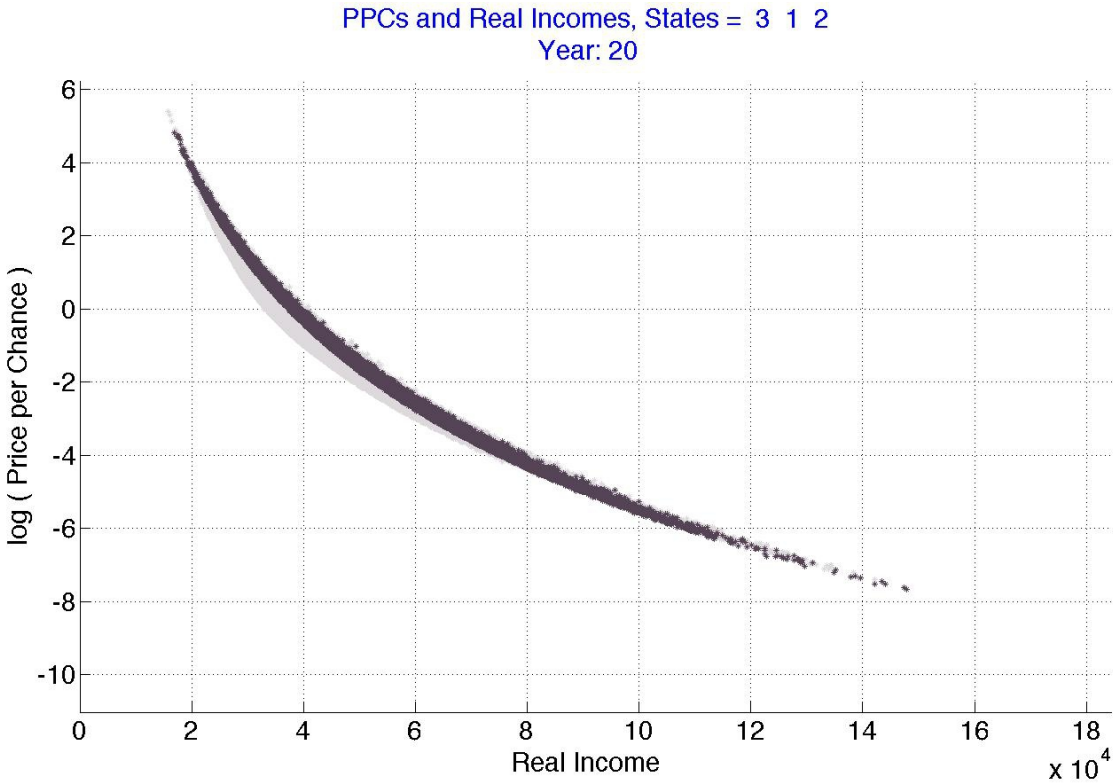

Fortunately, the inefficiency is very slight, as the yearly present value bar chart shows:

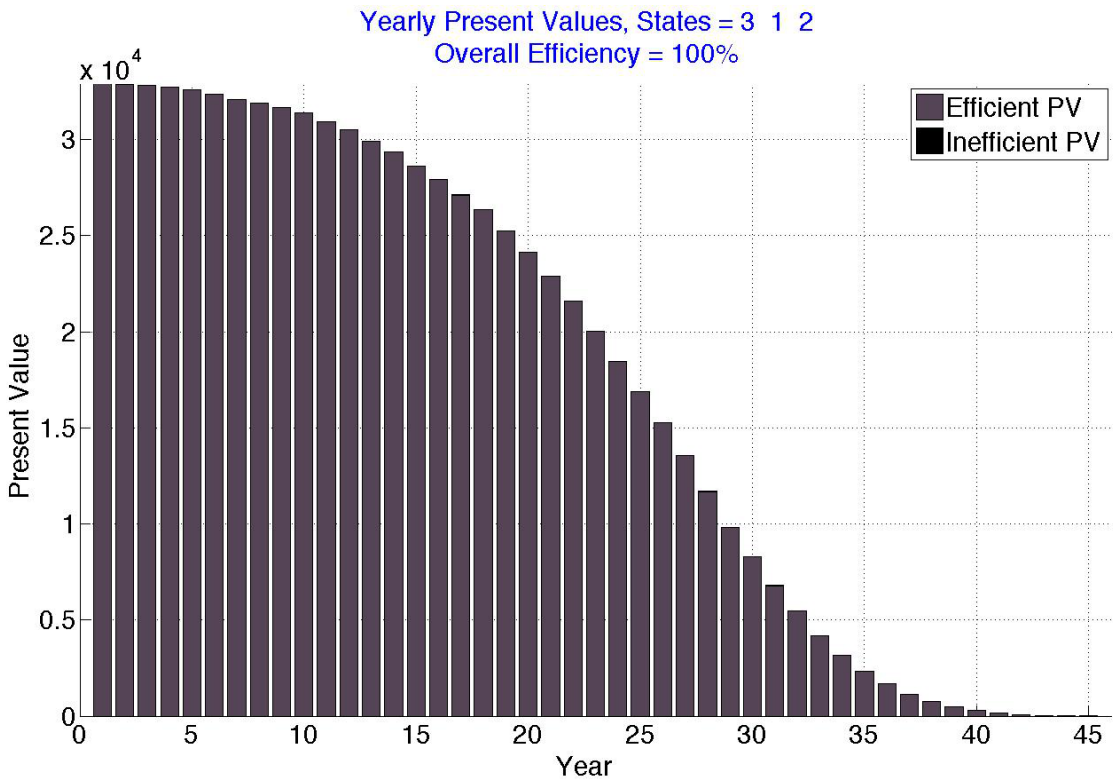

In no year is the cost efficiency sufficiently small to provide a visible black area at the top of a bar. This implies that the majority of points for each year in the prior PPC/Income graph are very close to a curve. The strategy undoubtedly costs a tiny bit more than need be, but the overall difference is sufficiently small that it doesn't show up in the graph since the overall efficiency percent measure is rounded to one decimal place. If the main results, including those in the prior three graphs are suitable for the retiree(s), only a purist would complain about such a minuscule amount of cost inefficiency.

Proportional Spending with a Decreasing Glide Path

Some proponents of proportional spending policies favor constant proportion strategies, such as rebalancing holdings annually to a 60%/40% stock/bond mix. Others advocate a *glide path,* with periodic rebalancing to a predetermined (and different) proportion mix of asset classes. Some choose to decrease portfolio risk over time, others to increase it. The number of possible combinations is large, providing fodder for new journal articles and financial product offerings.

We will be content to analyze two possibilities, using our market portfolio and TIPS investment vehicles. The first will move the asset allocation each year from 100% in the market portfolio to 0% in year 30 and thereafter. The second will begin with 0% in the market portfolio, moving to 100% in year 30 and thereafter. In each case the allocations will plot as linear functions.

We start by setting:

iPropSpending.glidePath = [1 0 ; 0 30];

Providing the desired glide path:

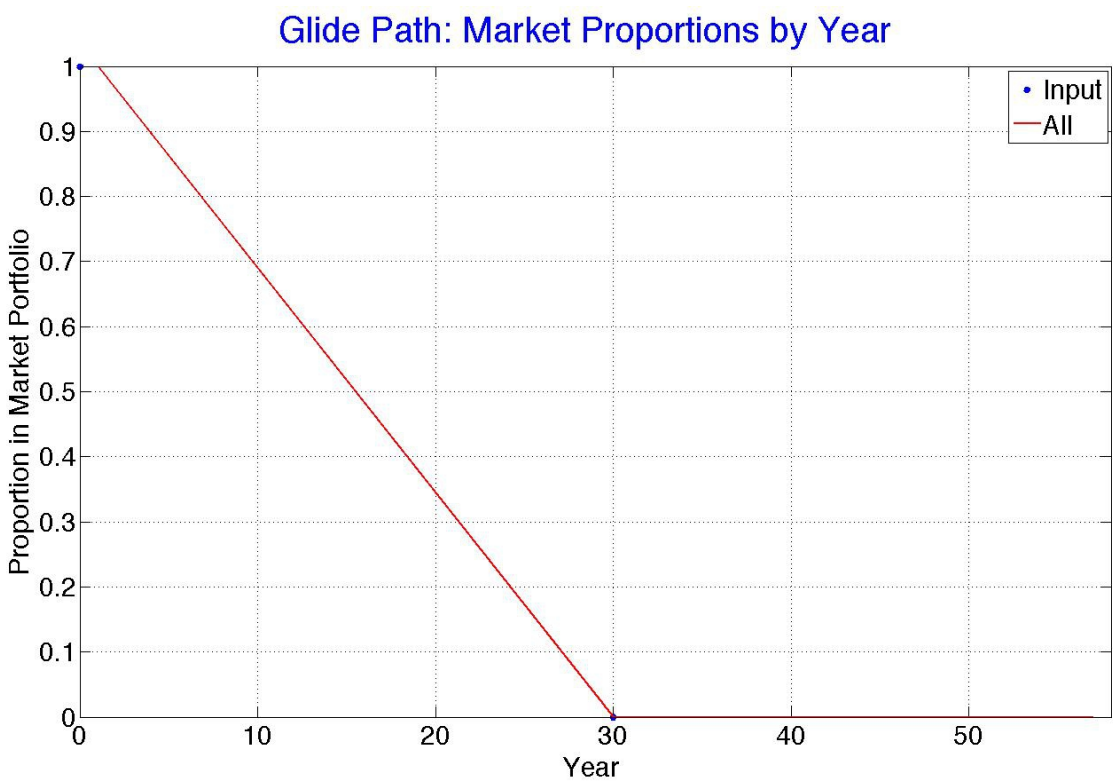

Not surprisingly, the ranges of incomes in each year differ from those in the previous case, with less variation in the later years. Here are the unconditional and conditional income maps:

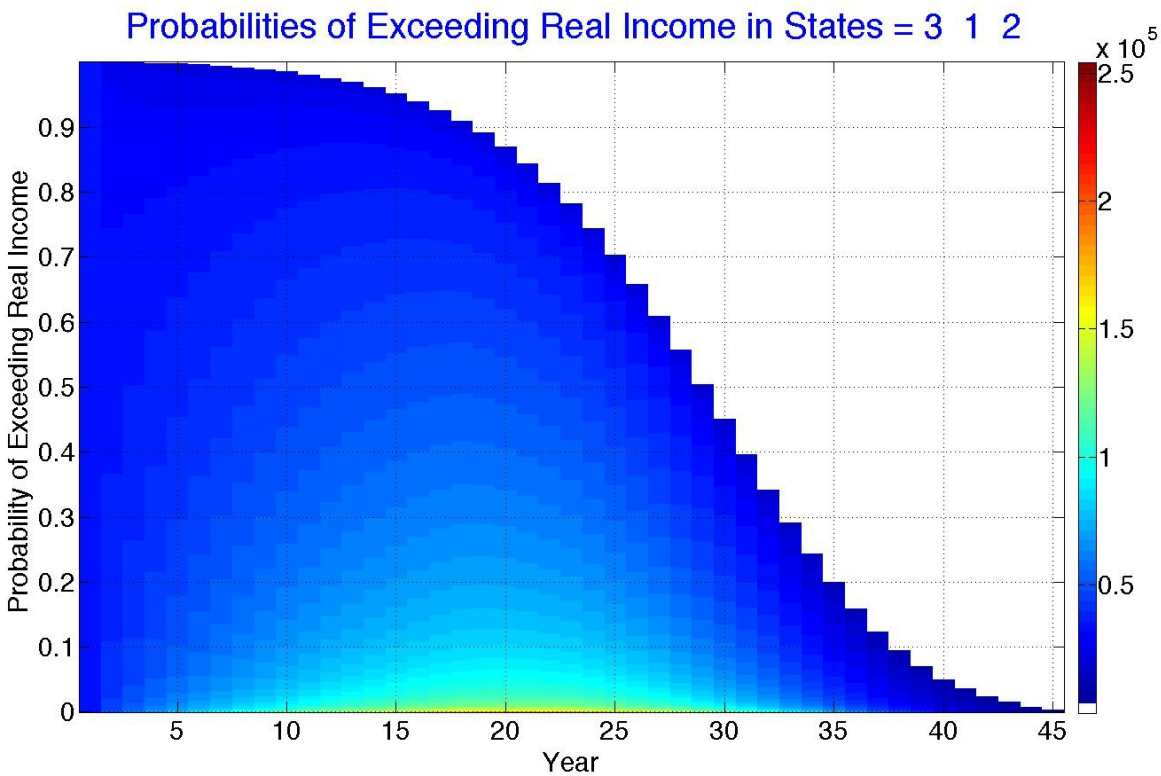

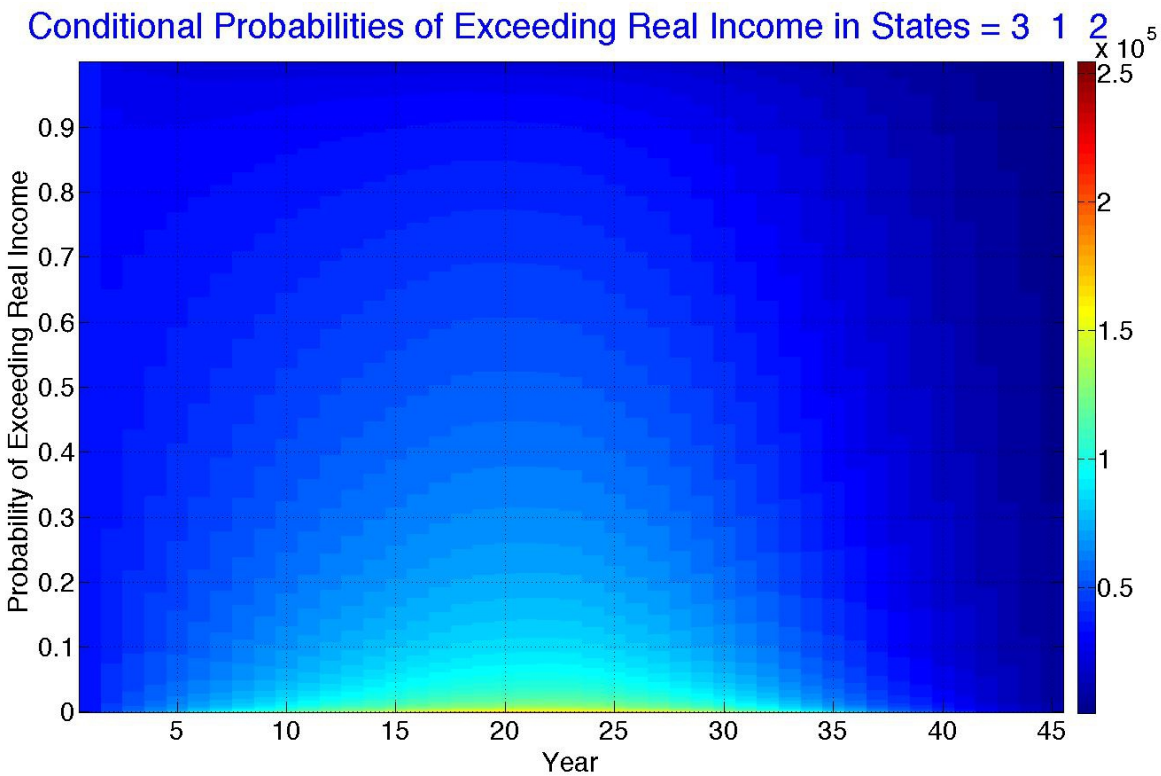

However, the relative allocations of present value among the possible recipients are almost the same as in each of the previous cases:

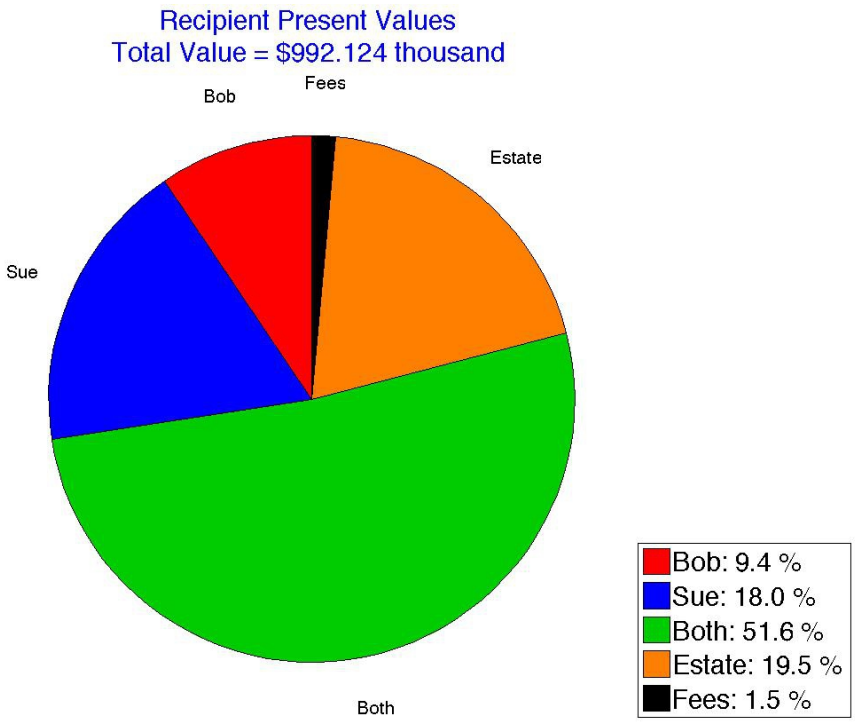

The graphs of yearly PPCS and real incomes differ, reflecting the different exposures to market risk over time. Moreover, in a given year there is more variation in income for a given PPC (scatter around the central tendency of the plot), reflecting greater cost inefficiency due to the path-dependence of cumulative portfolio return in each year. Here is the graph, stopped when showing results for year 20:

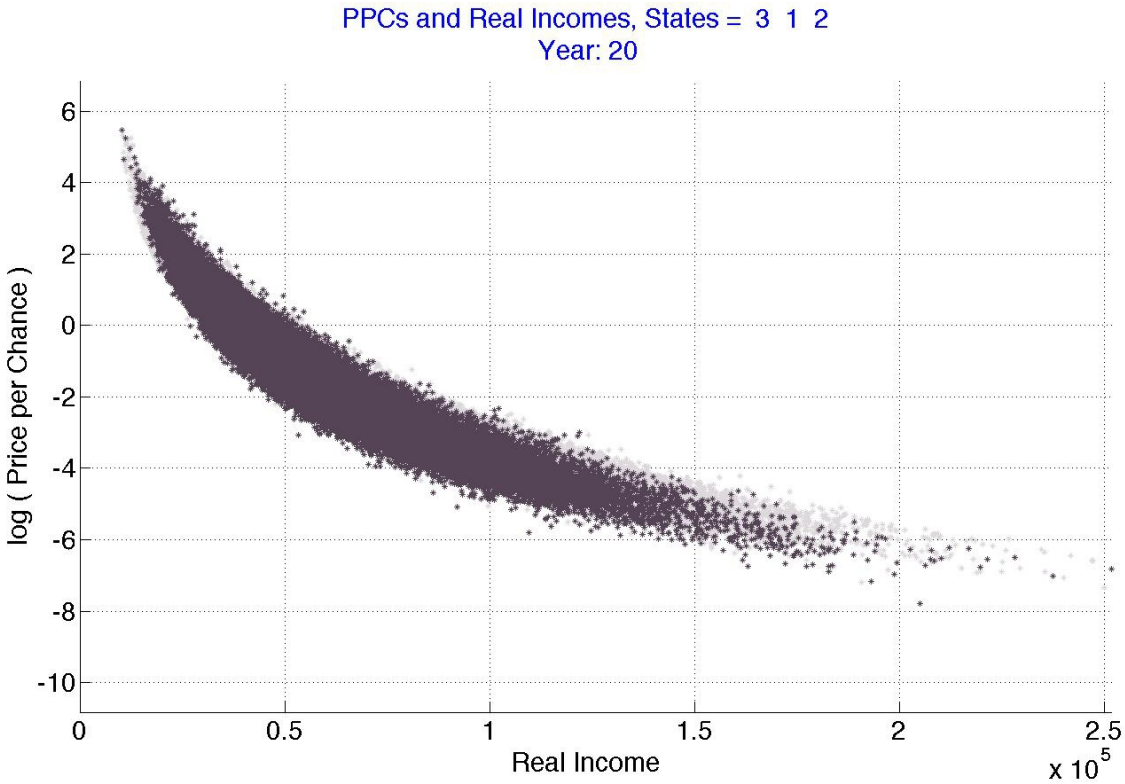

The inefficiency, while not trivial is not overwhelming. The entire set of probability distributions (one for each future year) could in principle be obtained for 98.4% of the amount invested ($984,000 instead of $1,000,000):

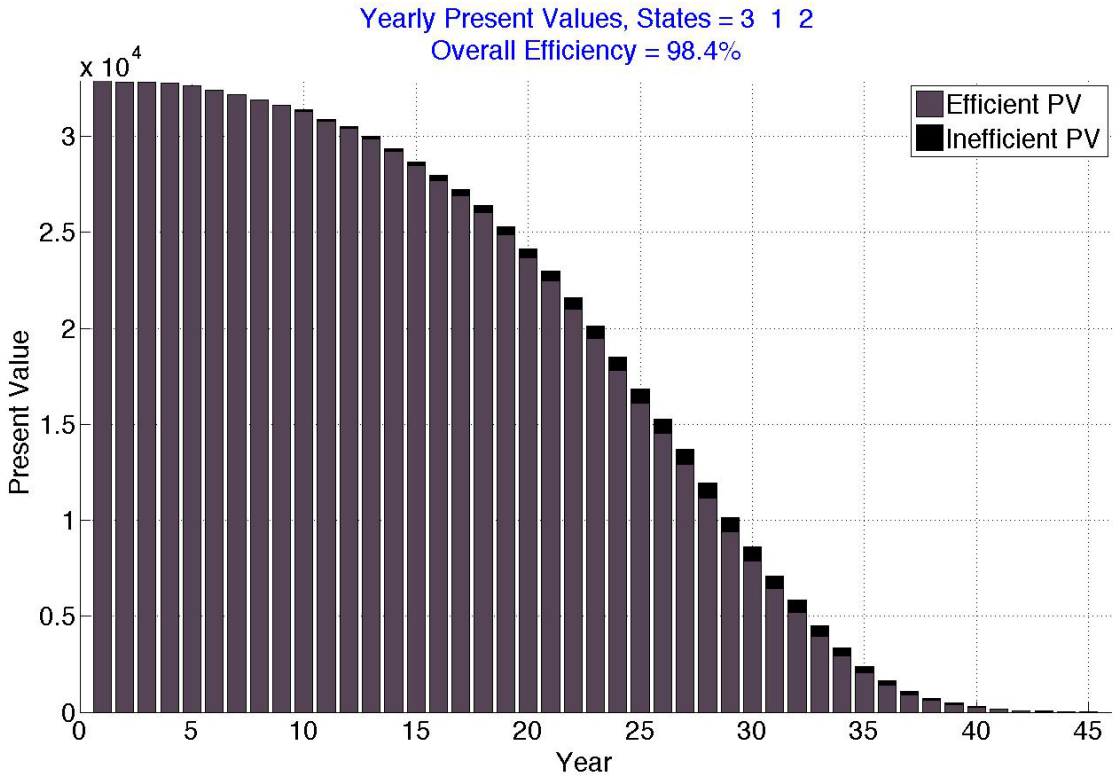

These graphs (and more) are available at:

http://www.stanford.edu/~wfsharpe/RISMAT/SmithCase_Chapter18.mp4

Proportional Spending with an Increasing Glide Path

For completeness, we conclude with an example in which the proportion invested in the market increases from year 1 through 30, then remains constant thereafter. We set the corresponding data element:

iPropSpending.glidePath = [0 1; 1 30];

Giving the following graph:

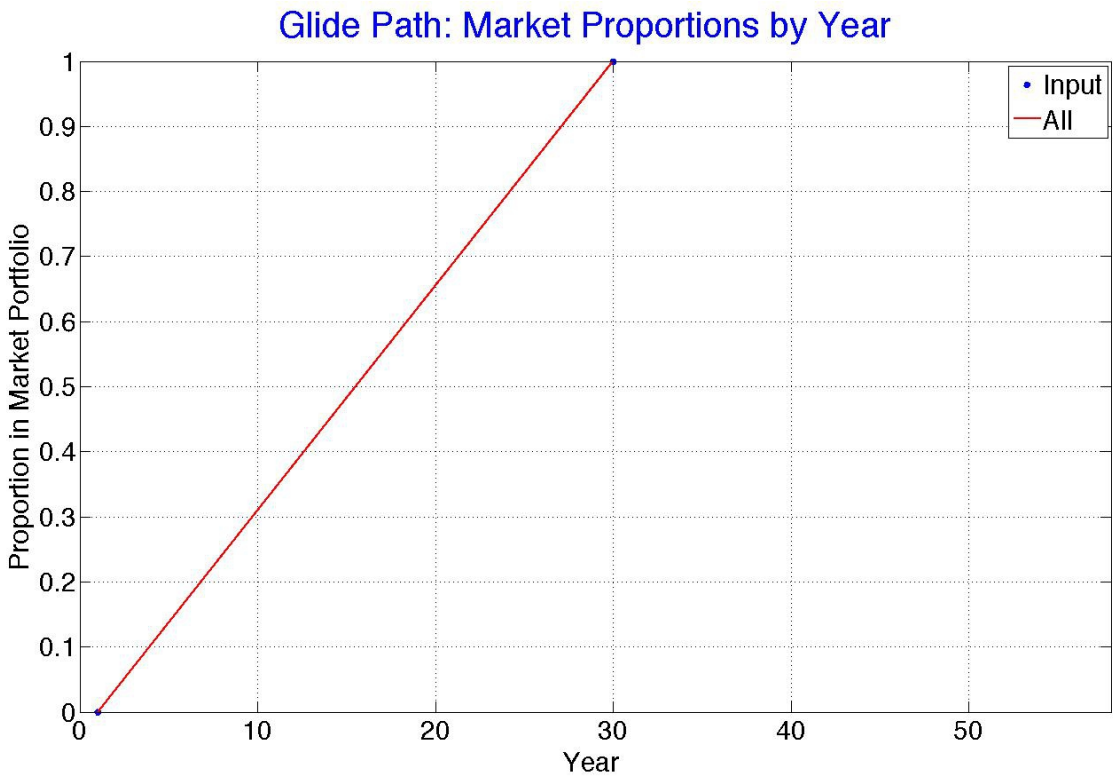

Since a picture is purported to be worth a thousand words, and this case is a variation on a theme we have pursued at length in this chapter, we present the next five graphs with minimal comments.

The changes in the income distributions from those for the decreasing glide path are, as one might expect, greater in the early years than in the latter ones:

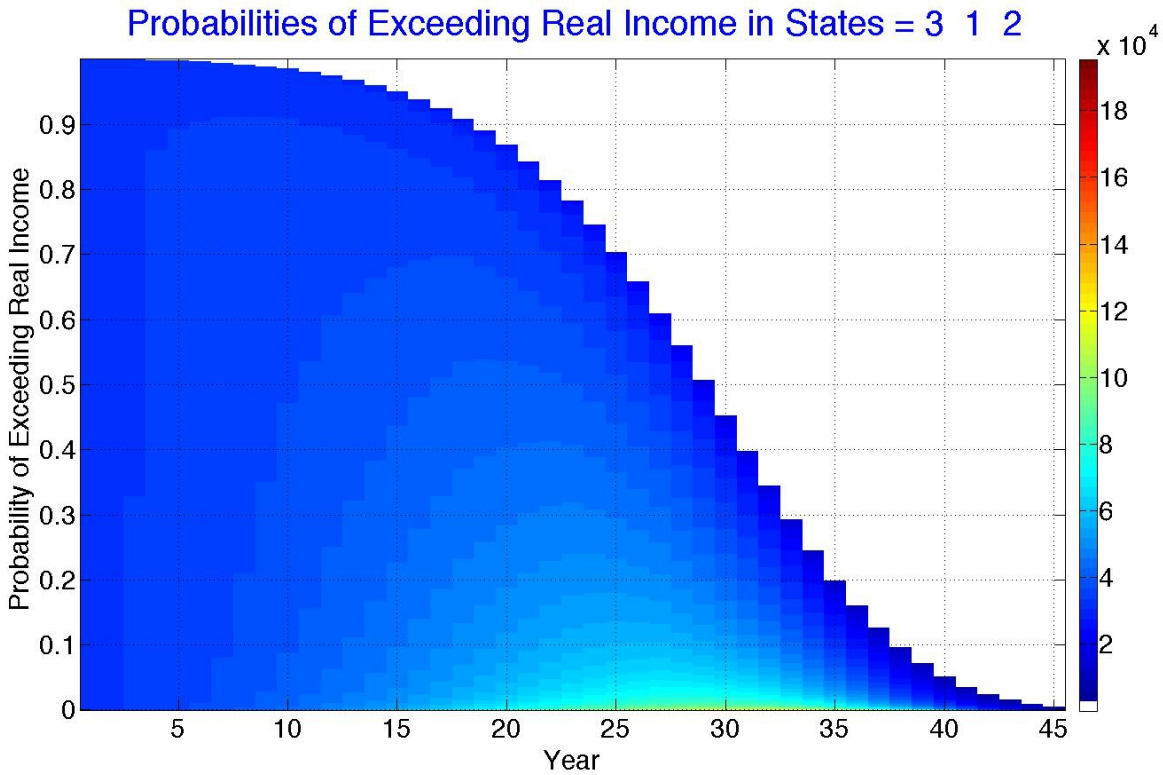

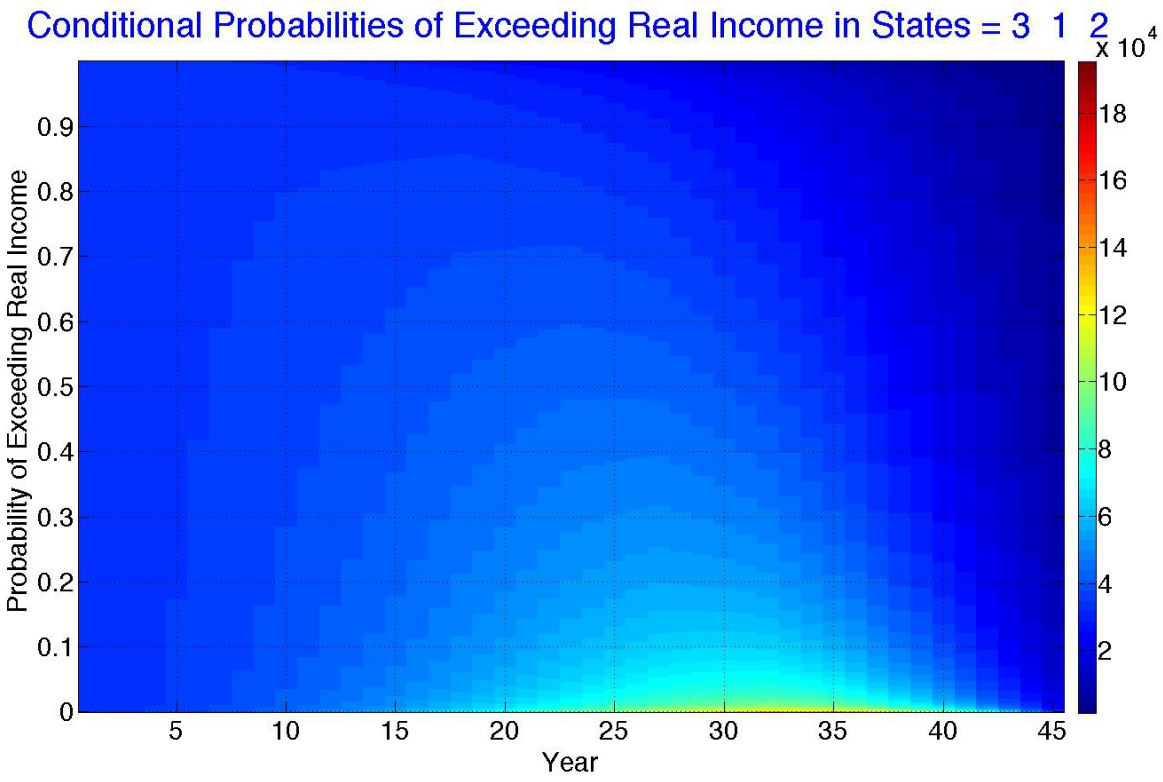
Conditional Probabilities of Exceeding Real Income in States = 3 1 2
x 10^4
Probability of Exceeding Real Income
Year
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
5
10
15
20
25
30
35
40
45
2
4
6
8
10
12
14
16
18

As we now know, any differences in the relative present values of the participants' claims for different investment strategies will be minor. Here are the results for this case:

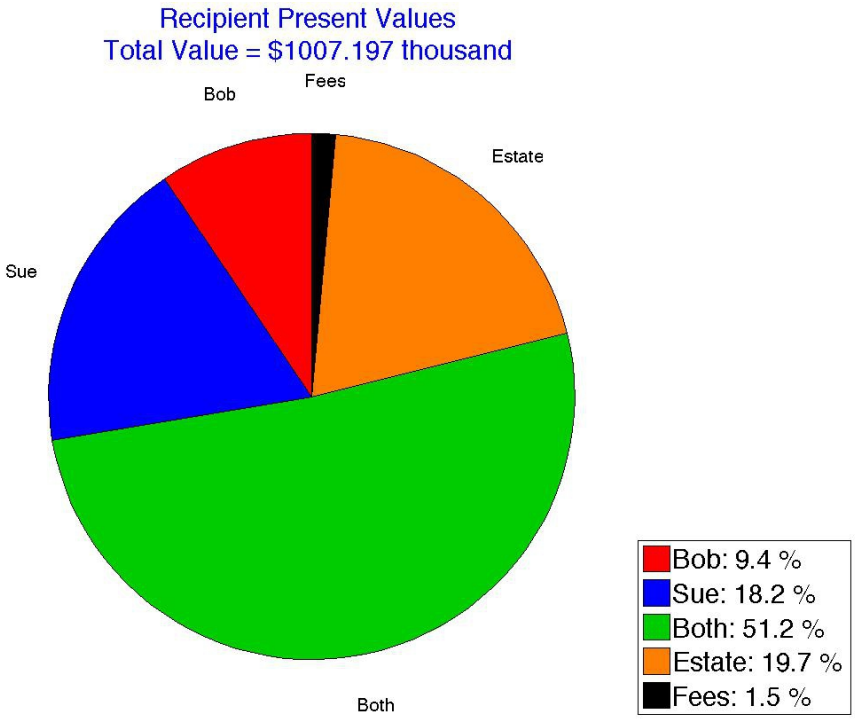

The implied marginal utilities will vary as one might expect, reflecting greater risk aversion in the early years and less in the later years. And there will be cost-inefficiency due to the path-dependency of the overall returns. That said, the cost efficiency is terribly low, although in this case it is slightly less than in the prior example.

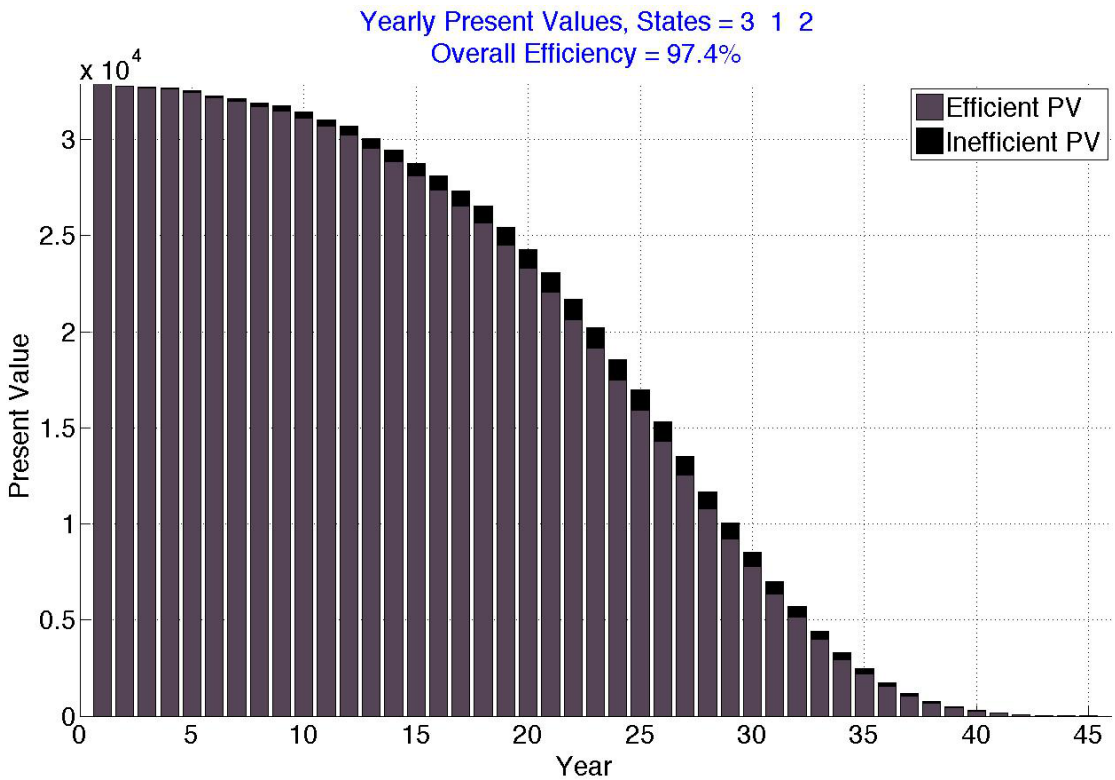

Proportional Spending and Lockboxes

We have repeated *ad nauseam* the equivalence of a traditional proportional spending policy to one that uses lockboxes,although in many cases the each of the lockboxes might have to contain shares in a fund that follows some sort of changing asset allocation policy, resulting in path-dependent and hence cost-inefficient income distributions. With this possibility included, any such proportional spending policy can be replicated with one that uses lockboxes. However, the converse does not follow. One might, for example, use only the market portfolio for the lockbox maturing in year 2 and only TIPS in the lockbox maturing in year 3. A traditional proportional spending policy cannot provide equivalent results.

For this and other reasons, for non-insured spending policies we prefer to focus on lockbox spending approaches, which are the subject of chapter 20. First, in Chapter 19 we examine *ratchets* – strategies designed to provide nominal incomes that may increase but will never (absent defaults) decrease.