

Retirement Income Analysis with scenario matrices

William F. Sharpe

17. Constant Spending

The 4% Rule

The October 1994 issue of the *Journal of Financial Planning* included an article by William P. Bengen titled “*Determining Withdrawal Rates Using Historical Data*” that has profoundly influenced financial practice ever since. To quote from the introduction:

At the onset of retirement, investment advisors make crucial recommendations to clients concerning asset allocation, as well as dollar amounts they can safely withdraw annually, so clients will not outlive their money. This article utilizes historical investment data as a rational basis for these recommendations ...

Citing a prior article by Larry Bierwirth, Bengen argues that “.. it pays to look .. at what actually has happened, year-by-year, to investment returns and inflation in the past”. He then proceeds to “... rely on actual historical performance of investments and inflation, as presented in Ibbotson Associates' *Stocks, Bonds, Bills and Inflation: 1992 Yearbook*.”

His initial analyses are based on the assumption that the client “.. continually rebalances a portfolio of 50 percent common stocks and 50 percent intermediate-term treasuries” but given the fact that the underlying database includes annual returns he presumably assumes annual rebalancing. He considers different asset allocations, concluding that “.. the 50/50 stock/bond mix appears to be near-optimum for generating the highest minimum portfolio longevity for any withdrawal scheme.” After further analysis he concludes that “Somewhere between 50-percent and 75-percent stocks will be a client's 'comfort zone”.

Bengen's analysis focuses on portfolio longevity: “how long the portfolio will last before all its investments have been exhausted by withdrawals.” He analyzes cases which assume a minimum requirement of 30 years of such longevity, concluding that “.. a first-year withdrawal of 4%, followed by inflation-adjusted withdrawals in subsequent years should be safe. In no past case has it caused a portfolio to be exhausted before 33 years, and in most cases it will lead to portfolio lives of 50 years or longer.” His conclusion is based on analyses of asset class returns for every 30-year period beginning in 1926 and ending in 1976 (for years after 1992 he assumes constant returns for each asset class using average returns from 1926 through 1992).

Importantly, Bengen advocates *constant spending* in real dollars every year, no matter how the underlying portfolio performs. Portfolio value at the outset is all-important, since the first year's spending is calculated as a specified percentage of that value. Thereafter the portfolio value has no effect on the real amount spent (unless it falls below one year's amount, in which case the entire remaining value is spent). Hence the title of this chapter.

In May 2012, in an online article under the title “*How Much is Enough?*” Bengen, referring to himself in 1994 as “... an obscure advisor in El Cajon, Calif.”, noted that after publishing the earlier paper, “.. the advisor increased his recommendation to a first-year withdrawal rate of 4.5%.” but noted that the conclusions of the original paper have been “... popularly enshrined, for better or worse, under the moniker of 'The 4% rule'.” Continuing to use actual historic periods with some splicing of divergent dates, he counsels: “I offer the following informal rule for your consideration: Take some pre-emptive action, no matter how mild, when the current withdrawal first exceeds the initial withdrawal rate by 25%.” Moreover, “I also offer the following corollary rule: If, despite initial action, withdrawal rates continue to rise, take more aggressive action... it's really tough to deal with double-digit withdrawal rates. Don't let them get that high.”

Since Bengen's initial paper, others have offered analyses of constant spending strategies using bootstrap procedures with simulated future asset returns are drawn randomly from actual historic annual returns, Monte Carlo simulations using returns drawn from probability distributions with assumed parameters, and so on. Not surprisingly, results have varied.

In the fall of 2016, a Business Week article appeared, titled “*Living on 4 Percent – Or Less*”. In it, Evan Inglis, an actuary at Nuveen Asset Management, is quoted as offering an alternative rule: “Divide your age by 20 – couples should use the younger partner's age – to get the percentage that you can safely spend.” For Bob and Sue, that would be 3.25% (65/20) of the initial portfolio value, followed by the same real income unless (until?) they run out of money.

In one form or another, constant real spending approaches remain ubiquitous. In 2016, Vanguard's web site offered a calculator for those asking “How much can I withdraw in retirement?” The inputs include:

Portfolio balance at retirement

Asset allocation (Conservative, Moderate, Aggressive)

Time spent in retirement

The outputs are:

Initial withdrawal rates

Initial monthly withdrawal

Here are the results for nine different sets of inputs:

	30 years	35 years	40 years
Conservative allocation	3.4 %	3.0 %	2.7 %
Moderate allocation	3.8 %	3.4 %	3.2 %
Aggressive allocation	4.0 %	3.6 %	3.4 %

The results assume:

“... you will increase the dollar amount of withdrawals you make over time to match the rate of inflation”

“... you will maintain a constant asset allocation over your entire planning horizon by rebalancing your portfolio at the end of each year

“... a conservative asset allocation is considered to be 20% stocks / 80% bonds; a moderate asset allocation is 50% stocks / 50% bonds; and an aggressive asset allocation is 80% stocks / 20% bonds.”

and:

“IMPORTANT:*The estimated withdrawal rate assumes an 85% chance that the portfolio won't run out of money before the end of your chosen investment horizon. The projections or other information generated by this tool are hypothetical, don't reflect actual investment results, and aren't guarantees of future results. Based on your input, results may vary with each use and over time.”*

Clearly, Monte Carlo analysis is employed:

“Return assumption. The tool uses forward-looking expectations for the U.S. and international capital markets generated by the Vanguard Capital Markets Model (VCMM). The VCMM is a proprietary financial simulation tool developed and maintained by Vanguard Investment Counseling & Research and the Investment Strategy Group. The VCMM uses a statistical analysis of historical data for interest rates, inflation, and other risk factors for global equities, fixed income, and commodity markets to generate forward-looking distributions of expected long-term returns. The asset return distributions are drawn from 10,000 simulations from the VCMM, reflecting 30 years of forward-looking simulations through December 2012.”

This is but one of many such analytic tools available online. In one form or another, constant spending approaches are alive and well.

As one might imagine, variants on the original rules abound in the financial planning literature. Some authors argue against using a constant asset allocation with annual rebalancing, recommending instead time-dependent or market-dependent allocations. Another school recommends following a glide path in which the proportions of asset classes are varied to decrease overall portfolio risk from year to year. Yet others advocate glide paths with increasing portfolio risk. But most continue to advocate constant real spending with little or no regard for variations in portfolio values.

Jason Scott, John Watson and I analyzed some of the characteristics of such approaches in two papers:

"Efficient Retirement Financial Strategies" in John Ameriks and Olivia Mitchell, *Recalibrating Retirement Spending and Saving*, Oxford University Press, 2008

"The 4% Rule -- At What Price?", *Journal of Investment Management*, Vol. 7, No. 3, Third Quarter 2009, pp. 31-48.

Pre-publication versions of each are available at stanford.edu/~wfsarpe under the heading *"Retirement Financial Strategies"*.

The abstract for the second paper provides a summary of our conclusions:

“The 4% rule is the advice most often given to retirees for managing spending and investing. This rule and its variants finance a constant, non-volatile spending plan using a risky, volatile investment strategy. As a result, retirees accumulate unspent surpluses when markets outperform and face spending shortfalls when markets underperform. The previous work on this subject has focused on the probability of shortfalls and optimal portfolio mixes. We will focus on the rule’s inefficiencies—the price paid for funding its unspent surpluses and the overpayments made to purchase its spending policy. We show that a typical rule allocates 10%-20% of a retiree’s initial wealth to surpluses and an additional 2%-4% to overpayments. Further, we argue that even if retirees were to recoup these costs, the 4% rule’s spending plan often remains wasteful, since many retirees may prefer a different, cheaper spending plan.”

Speaking for myself, and putting the subject in non-academic terms: it just seems silly to “finance a constant, non-volatile spending plan using a risky, volatile investment strategy.” One would hope that prudent financial advisors who adopt such calculations would at least counsel annual reviews in which planned real spending for the forthcoming year could be revised to take into account portfolio performance over the prior twelve months. Perhaps advisors who employ the apparatus used to derive the initial constant real spending could do so anew each year, thereby providing more sensible advice (and, of course, earning additional fees).

Not surprisingly, here we advocate different approaches (to be developed in subsequent chapters). But for completeness, this chapter will provide and illustrate the use of functions for constant real spending strategies.

The iConstSpending Data Structure

We start, as is our custom, with a function for the creation of a data structure for a constant spending source of income:

```
function iConstSpending = iConstSpending_create( )
    % create a constant spending income data structure

    % amount invested
    iConstSpending.investedAmount = 100000;

    % proportion of initial investment spent in first year
    iConstSpending.initialProportionSpent = 0.040;

    % relative incomes for personal states 1,2 and 3
    % (any remaining value is paid in personal state 4)
    iConstSpending.pStateRelativeIncomes = [ 0.5 0.5 1.0 ];

    % graduation ratio of each real income relative to the prior income distribution
    iConstSpending.graduationRatio = 1.00;

    % matrix of points on 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
    iConstSpending.glidePath = [ 1.0 ; 1 ];

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

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

end
```

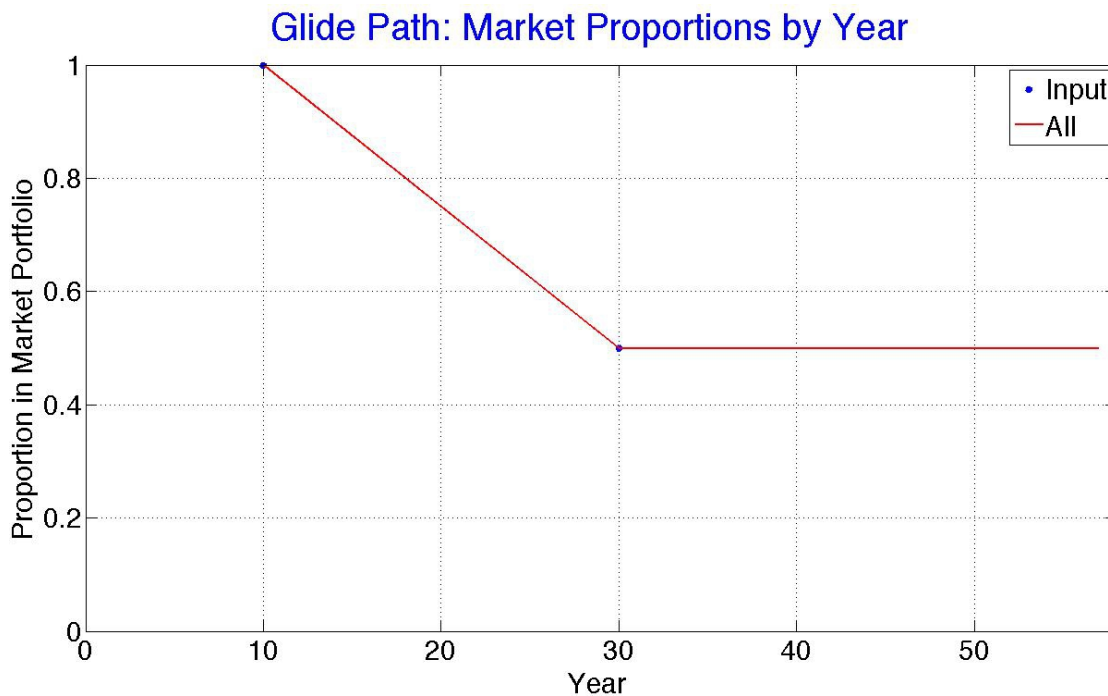
The first elements are relatively straightforward. The amount to be invested is stated in dollars and the proportion to be used as income for the first year as a ratio (here, 0.040 for the 4% rule). Although rarely used in this context, for generality we include our usual options for changing the amounts to be paid based on personal states (here, half as much when only Bob or Sue is alive), and a graduation ratio to adjust the real amount spent each year by some factor (we set the initial value to 1.0 so there will be no such adjustment).

The next element allows for the creation of a glide path for the fund's asset allocation. This is described with a matrix of points on a graph with the year on the horizontal (x) axis and the proportion to be invested in the market portfolio in that year on the vertical (y) axis (the proportion to be invested in TIPS is equal to 1 minus the proportion invested in the market portfolio). The market proportions should be in the top row of the matrix, with the corresponding years in the bottom row. When the matrix is processed, a proportion (y) is computed for every year (x). For years prior to the first listed, the proportion is the first one shown; for years after the last listed, the proportion is the last one shown. For each of the years listed in the second row, the proportion is the corresponding amount in the first row. The proportion for each of the remaining years is computed using a linear function for the two adjacent years shown.

For example, consider a case in which:

`iConstSpending.glidePath = [1.0 0.5 ; 10 30];`

The resulting glide path, produced by setting `iConstSpending.showGlidePath = 'y'`, is shown below (but the horizontal red line for years 1 through 10 showing the proportion equal to 1.0 is hidden by the grid line in this case):



Our approach to defining a glide path makes it possible to set any desired pattern of asset proportions over future years. We will also use this construction in the next chapter. However, as we will see, there are better ways than using a glide path to change probability distributions of future income.

The last data element indicates the retention ratio – the proportion of year-end portfolio value remaining after the investment manager or managers have taken their portion. The default value is 0.999, indicating a total expense ratio of 0.001% (10 basis points per year). As discussed in earlier chapters, while index funds may take this amount or less, for actively-managed and specialty funds the retention ratio may be as low as 0.990 or worse. And as the software will indicate, such expense ratios may divert considerable amounts from one's saving to pay fees to investment managers.

Processing an *iConstSpending* Data Structure

Given the many options we have included, processing the *iConstSpending* data structure requires some rather tedious operations. As usual, the reader is invited to consider skimming or skipping the details. For those with a deeper interest, we take sections of the function in turn.

Here is the first set of statements:

```
function client = iConstSpending_process( iConstSpending, client, market );

    % get matrix dimensions
    [nscen nyrs] = size( market.rmsM );
    % get glidepath
    path = iConstSpending.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;
```

This is mostly housekeeping, designed to adjust the glide path inputs so that the proportions range between 0 and 1 inclusive, append points for the first and last year if needed, and insure that all the points are in the order of increasing income (x) values. Boring housekeeping, to be sure, but useful.

The next section creates vectors with the coordinates of points for each year, with the x-values (years) in a *pathxs* vector and the y-values (market proportions) in a *pathys* vector.

```
% 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) ) ];
```

This accomplished, the results can be shown in a graph such as the one included previously, if requested. The statements are:

```
% show glide path if desired
if lower( iConstSpending.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;
end; % if lower(iConstSpending.showGlidePath) == 'y'
```

Next, we create a complete matrix with gross returns for the investment policy:

```
% 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;
```

The next statements get the retention ratio, set a vector of portfolio values to the initial amount, then create a matrix with desired spending amounts, taking into account the graduation ratio and desired relative incomes in the personal states 1, 2 and 3:

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

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

% initialize desired spending matrix
desiredSpendingM = zeros( nscen, nyrs );
% create matrix of desired real spending for highest personal state
prop = iConstSpending.initialProportionSpent;
amt = prop * iConstSpending.investedAmount;
gradRatio = iConstSpending.graduationRatio;
factors = gradRatio .^ ( 0: 1:nyrs-1 );
% create matrix of maximum desired spending
maxSpendingM = ones( nscen, 1 ) * ( amt*factors );

% add amounts to desired spending matrix
props = iConstSpending.pStateRelativeIncomes;
props = props / max( props );
props = max( props, 0 );
for ps = 1:1:3
    s = maxSpendingM .* props( ps );
    m = ( client.pStatesM == ps ) .* s;
    desiredSpendingM = desiredSpendingM + m;
end;
```

Finally (!) we move through the matrices, year by year, computing the portfolio returns, computing and deducting from portfolio values the fees paid and incomes paid out at the beginning of each year, and adding incomes and fees to their respective matrices.

```
% 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 ) = 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 .* min( desiredSpendingM( :, yr ), portvals );
    % 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;
```

This tedious work finished, the function adds the incomes and fees for this strategy to the corresponding client matrices and graciously exits.

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

end
```

Analysis

Now to substance. We begin with cases in which the portfolio is invested entirely in the market in every year. This can be simple accomplished by setting:

iConstSpending.glidePath = [1.0 ; 1.0];

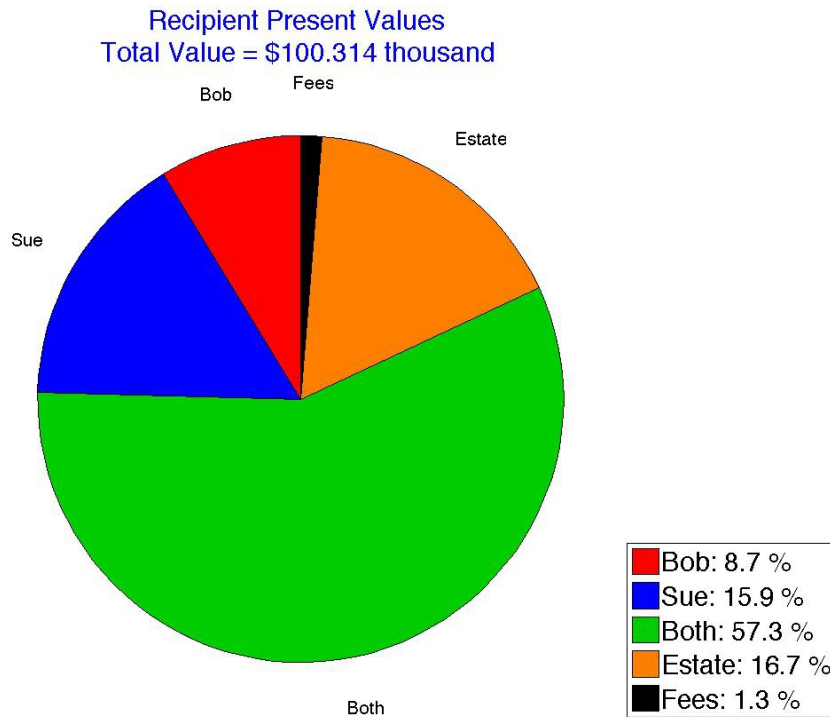
Note that this corresponds roughly with Bengen's original assumption that the 50% of the portfolio is invested in a stock portfolio and 50% in intermediate bonds since our market portfolio includes both bonds and stocks (in market proportions).

Also, to replicate other aspects of the original case we set:

iConstSpending.retentionRatio = 0.999;
iConstSpending.initialProportionSpent = 0.040;
iConstSpending.graduationRatio = 1.00;
iConstSpending.pStateRelativeIncomes = [1 1 1];

We include only expenses for a relatively low-cost index fund (retention ratio = 0.999) and choose to pay out a real amount each year equal to the legendary 4% of the initial portfolio value, avoiding any graduation in payments year to year or differences in amounts paid depending on which of our protagonists is extant. And, of course, any money left afterwards will go to the estate.

Here are the present values of the amounts paid:



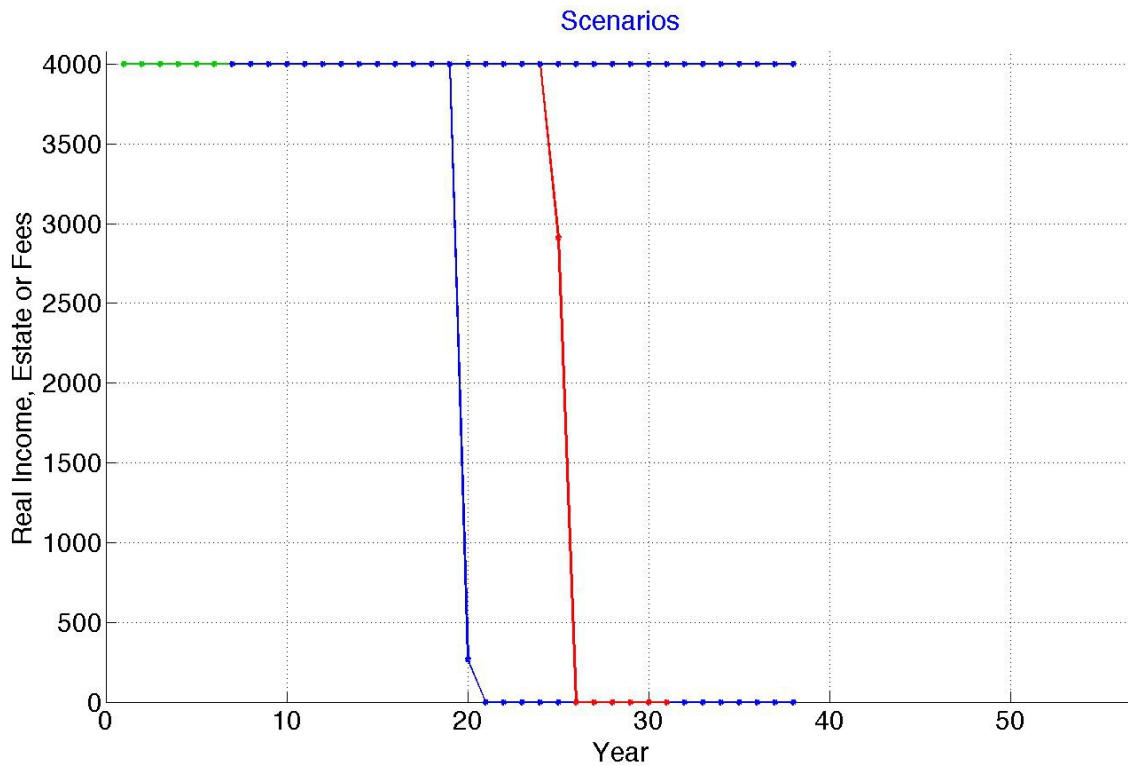
As we know, the total may not equal exactly the amount invested. In this case the present value of all the payments is \$100.314 thousand, while the actual amount invested was \$100.0 thousand. This is not something for nothing – simply the result of sampling error, as discussed in earlier chapters.

The striking aspect of this graph is the present value of the amounts that may go to Bob and Sue's estate (16.9% of the total value). In this sense a 4% payout leaves considerable money on the table. A few runs of the program with different payouts shows a likely range of corresponding results:

Initial Payout Ratio	Present Value of Estate (%)
3.0 %	30.3 %
3.5 %	22.8 %
4.0 %	16.7 %
4.5 %	12.4 %
5.0 %	8.9 %

The lower the payout ratio, the happier will be Bob and Sue's heirs. But their gain is Bob and/or Sue's loss.

The following figure shows 20 scenarios for a 4.0% payout in the first year:

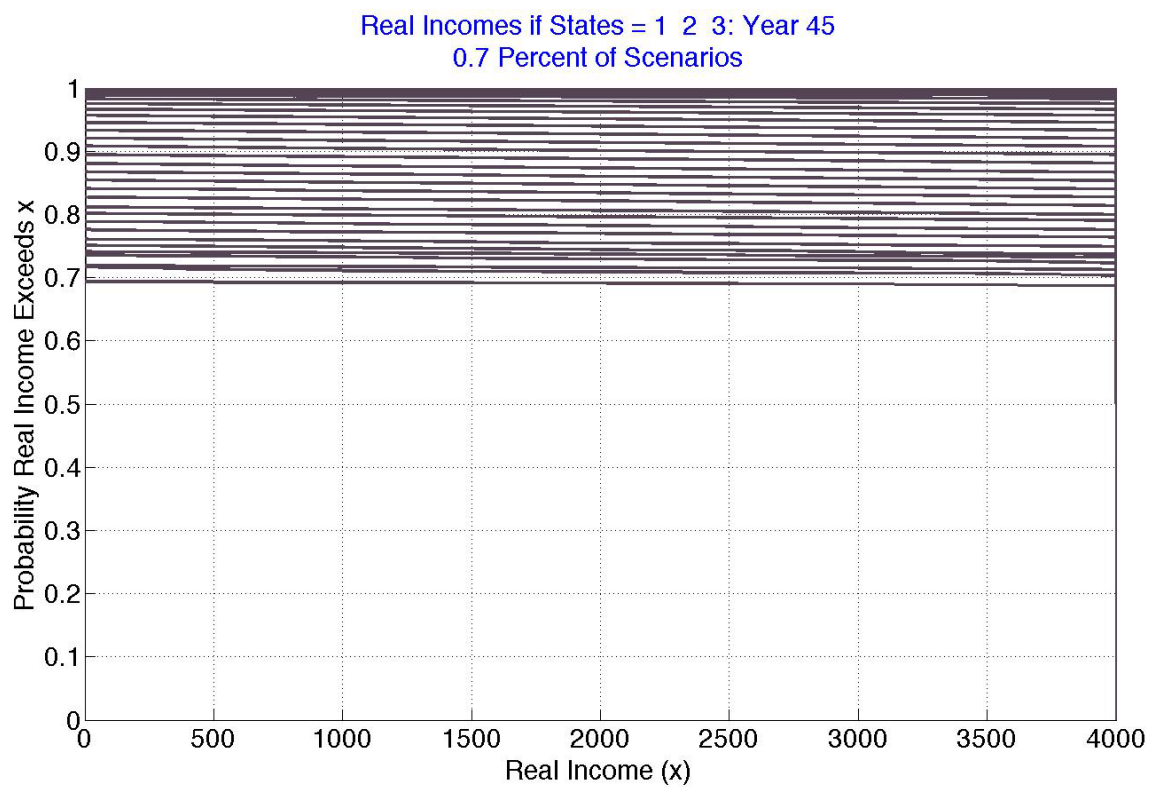


In one of the scenarios, Sue (shown in blue) lived long enough to run out of money, spending the remaining amount (roughly \$300) in year 20, then receiving no income from her savings thereafter. In another scenario, Bob (shown in red) survived Sue, receiving slightly less than \$3,000 from his savings in year 26, and nothing thereafter. This is the bad news. The good news is that in 18 out of 20 scenarios, Bob and/or Sue received the full amount of real income that the strategy was designed to provide.

The income distributions graph provides summary information across all scenarios in which Bob and/or Sue are alive. We choose the conditional version and include personal states 1,2 and 3:

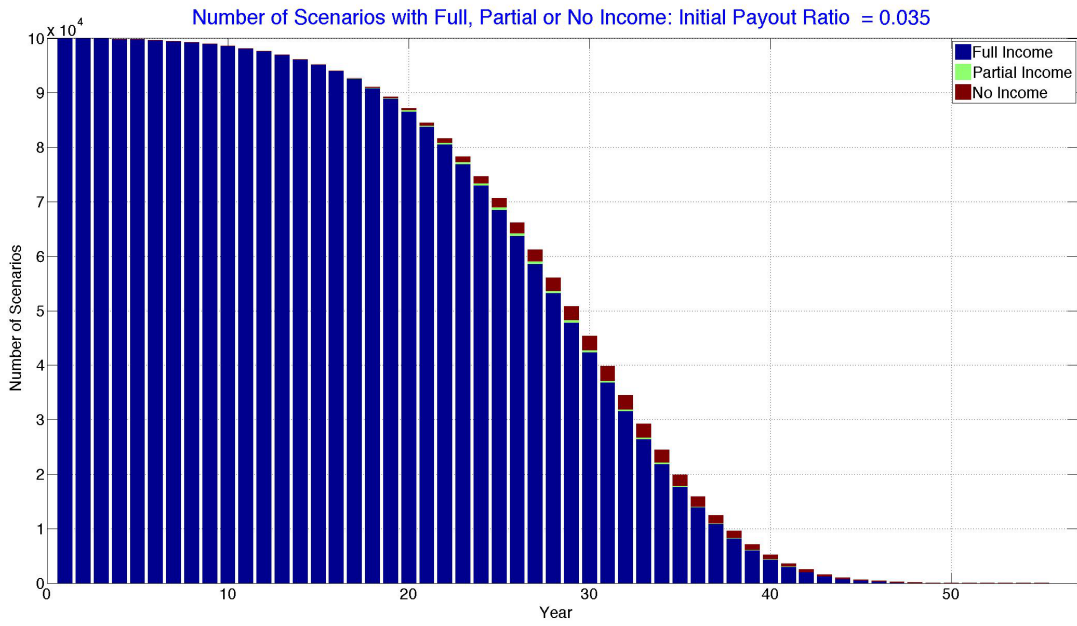
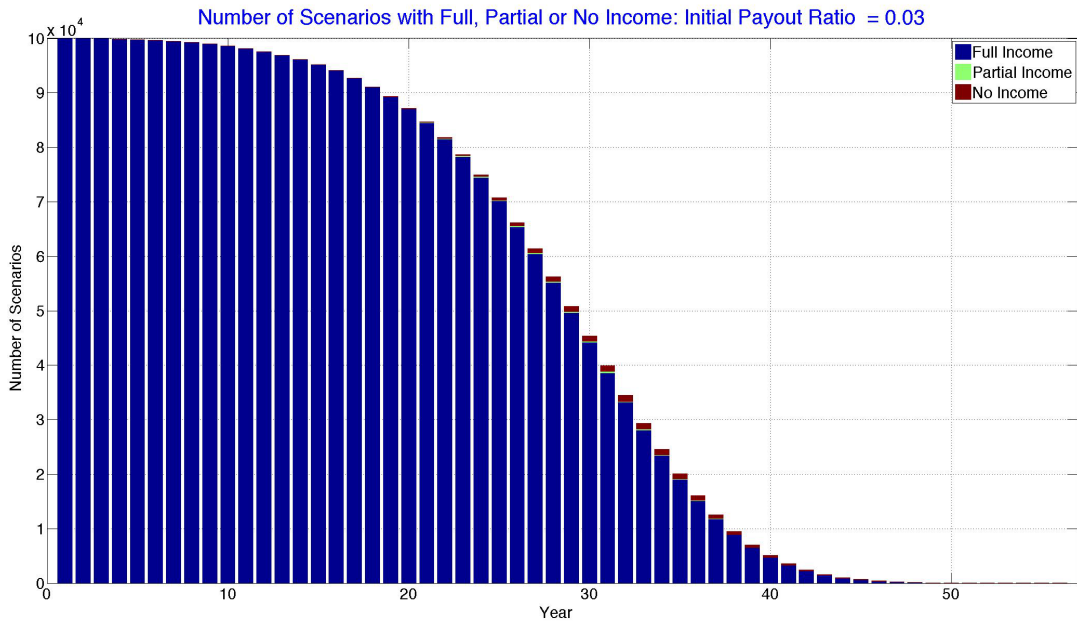
```
analysis.plotIncomeDistributions = 'y';  
analysis.plotIncomeDistributionsTypes = {'rc'};  
analysis.plotIncomeDistributionsStates = { [1 2 3] };
```

Here is the final version:

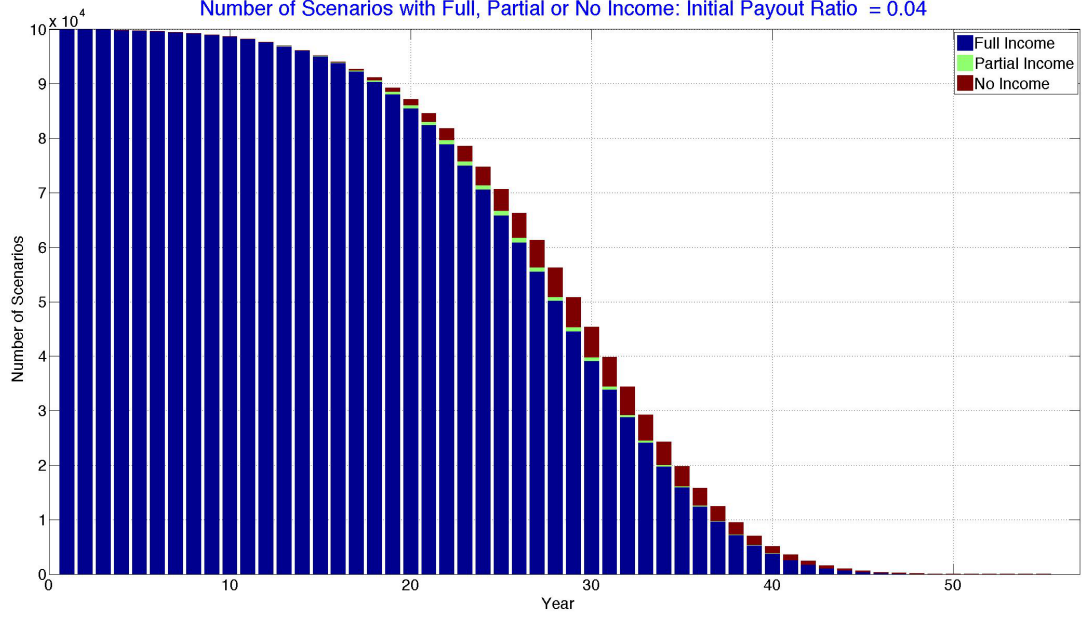


The curves for the first years are vertical lines at the promised income (\$4,000). But for later years they are almost flat step functions. The last one, for year 45, shows that in roughly 70% of the scenarios, real income was the full \$4,000, in a few it was slightly less, and in most of the rest it was zero. But, as shown, in only 0.7% of the scenarios was anyone alive in year 45, so this may not be devastating news.

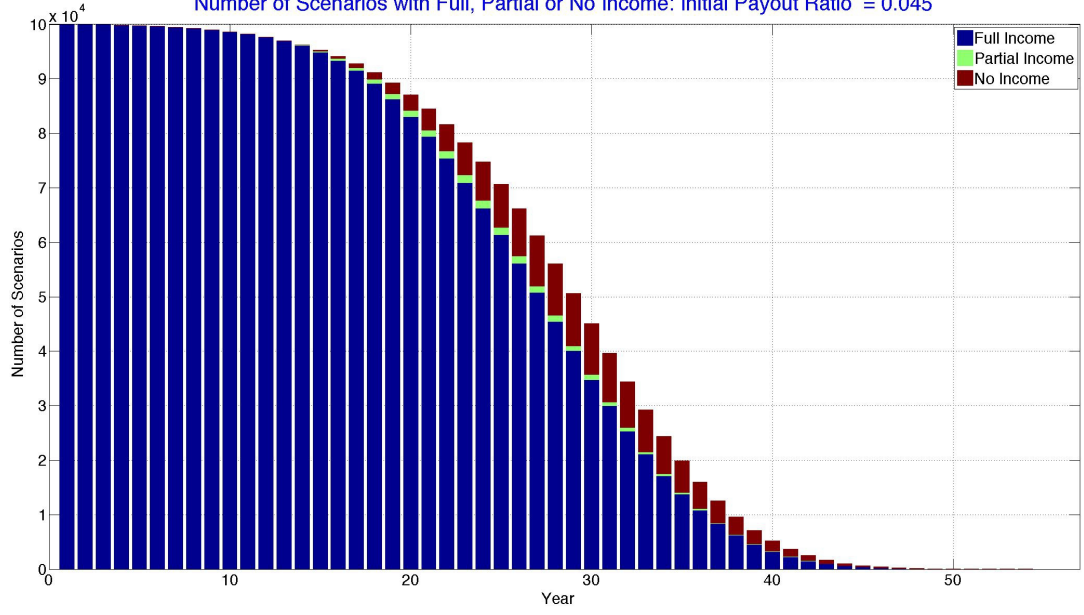
To get a better idea of the prospects of running out of money with a constant spending policy, one can count the number of scenarios in each year in which someone is alive (personal states 1, 2 or 3) and either (a) the full income is received, (b) partial income between zero and the full income is received and (c) no income is received. The following graphs show the results obtained with initial payout ratios of 3.0%, 3.5%, 4.0%, 4.5% and 5.0%.

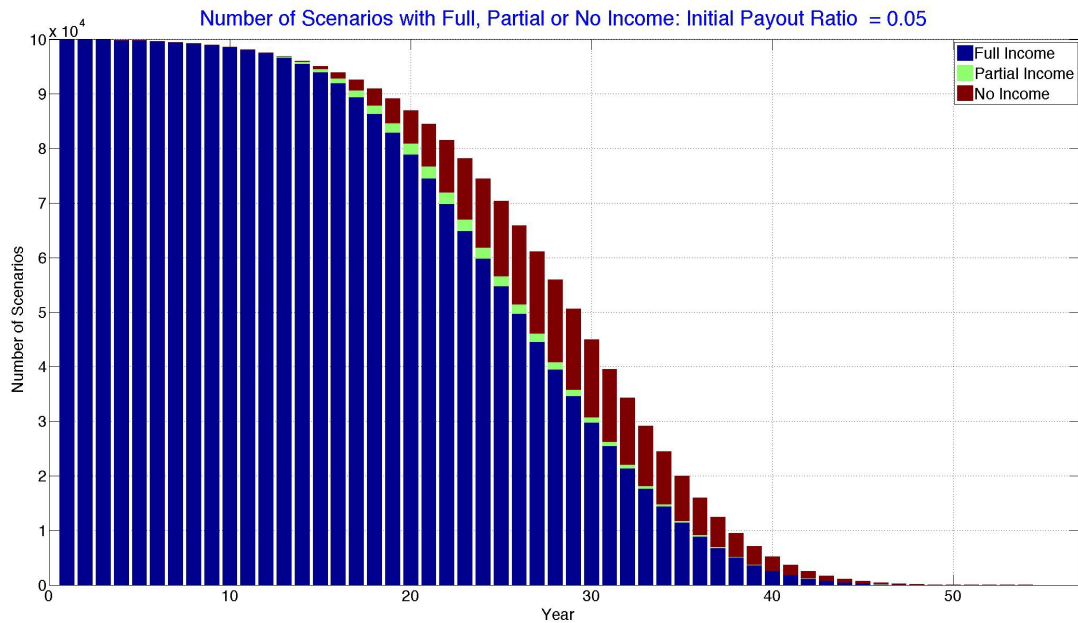


Number of Scenarios with Full, Partial or No Income: Initial Payout Ratio = 0.04



Number of Scenarios with Full, Partial or No Income: Initial Payout Ratio = 0.045

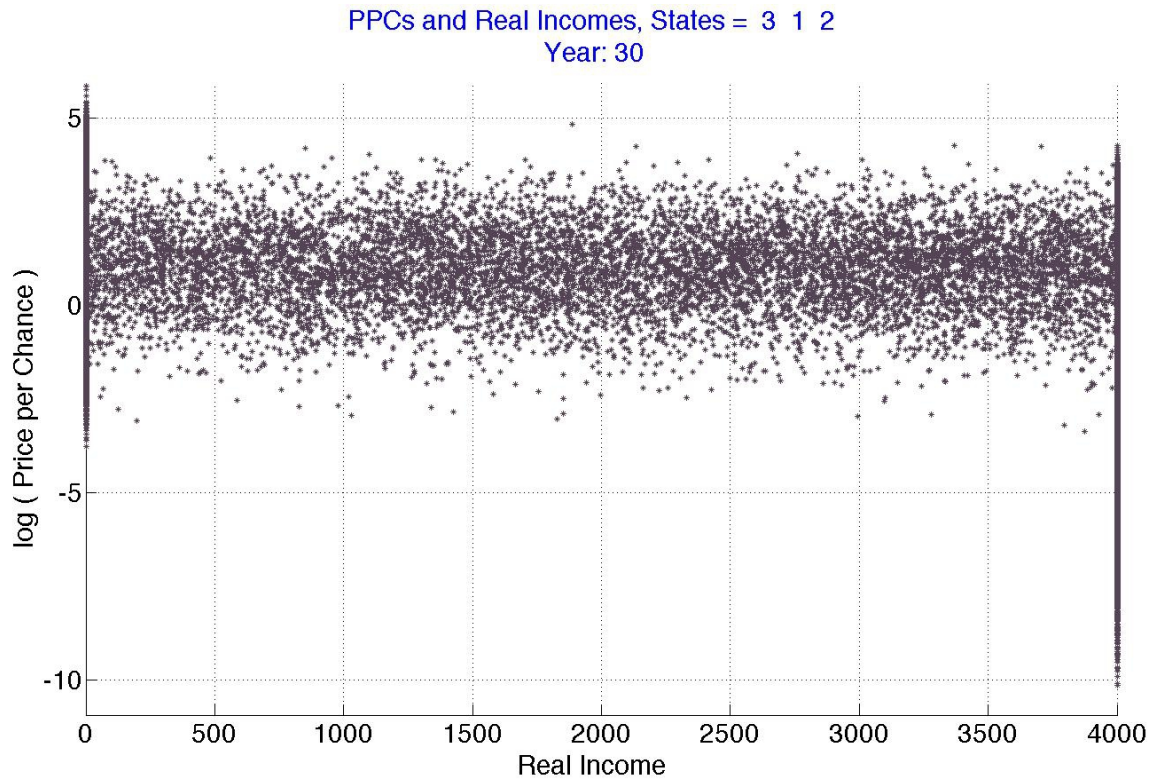




These graphs show the other side of the trade-off between the possible amounts received by Bob and Sue and those received by their estate. As we saw earlier, the higher the initial proportion spent, the larger the value of potential payments to Bob and Sue. But the higher the initial proportion, the greater also is the chance that they will run out of money. Consider an extreme case: if Bob and Sue want to spend all their savings while they are alive, the safest approach would be to spend everything in the first year. At the other extreme, they could take out only a tiny percentage of the initial savings, thereby virtually insuring that they would not run out of money, but living in penury and almost certainly leaving a large estate.

This sort of trade-off is not unique to constant spending approaches. Absent some sort of pooling of mortality risk, reserving some savings for payments in old age will lower incomes in earlier years and increase the probability of leaving a significant portion of savings to an estate. As we will see, this may argue for adopting an approach that calls for spending some portion of savings over a given number of future years, then using the remainder of savings to either purchase a deferred annuity at the present time or reserving it to purchase an immediate annuity at a later date. We will have more to say about this in later chapters.

Returning to the case at hand, it is interesting to see the relationship between PPC values and the chosen amounts of income. Here is the graph with results for the first thirty years for payments equalling 4.0% of an initial amount of \$100,000:

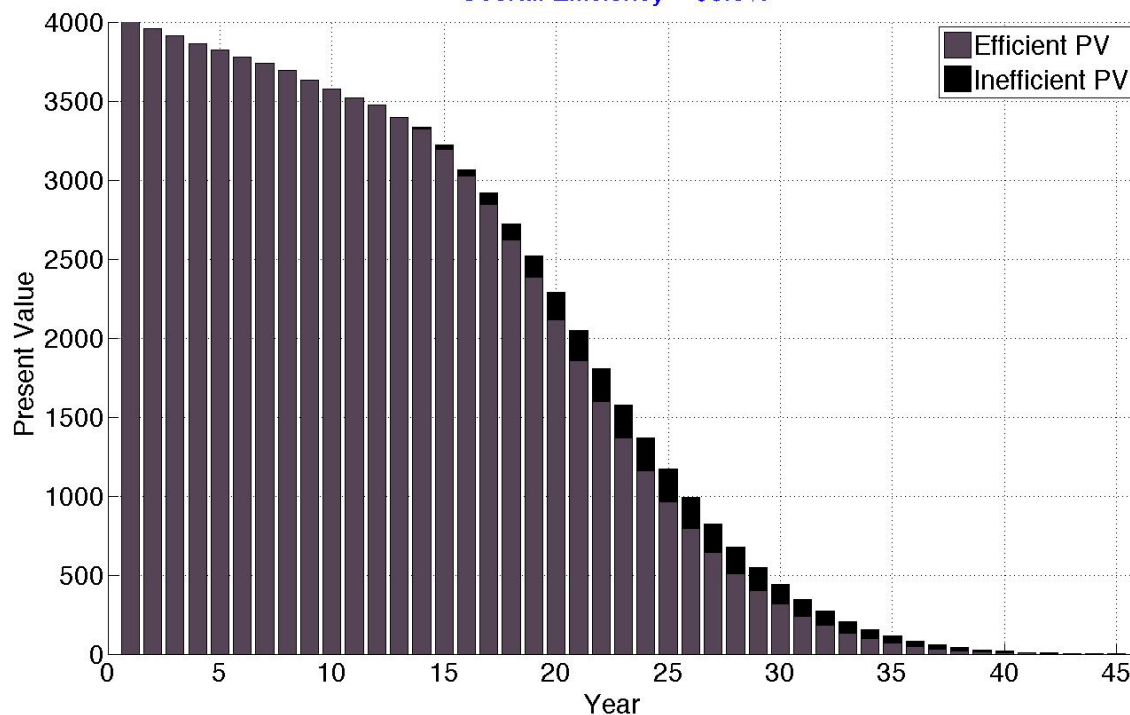


The majority of points fall at the maximum income of \$4,000 or the minimum of \$0, with those representing partial payments in between. There is some tendency for the maximum payments to be associated with lower PPC values and the minimum (zero) with higher PPC values. But this is by no means always the case. Moreover, the partial payments vary with no apparent relationship between PPC and real income.

The implication of these results is that constant spending rules are not efficient in a least-cost sense: it should be possible to obtain the same annual probability distributions of income at lower cost by obtaining higher incomes in states with lower prices (PPC values).

The yearly present value graph shows that this is the case; it also provides a measure of the overall cost efficiency:

Yearly Present Values, States = 1 2 3
Overall Efficiency = 96.5%



In this case, one could obtain the same probability distribution of real income in each year as that provided by the constant spending rule for a total cost equal to 96.5% of that in the example (\$96,500 instead of \$100,000). Why? Because the varying amounts received in each of the later years are not inversely related to prices (because they are not positively related to cumulative returns on the market portfolio). Even though in this case all the funds were invested in the market portfolio, the periodic withdrawals of fixed real amounts regardless of the performance of the portfolio caused later values to be dependent on the path of market returns rather than solely the cumulative return to the date of withdrawal. And, since in our equilibrium model, prices (PPC values) depend solely on the cumulative return on the market portfolio, the strategy is not cost-efficient.

Investment Glide Paths

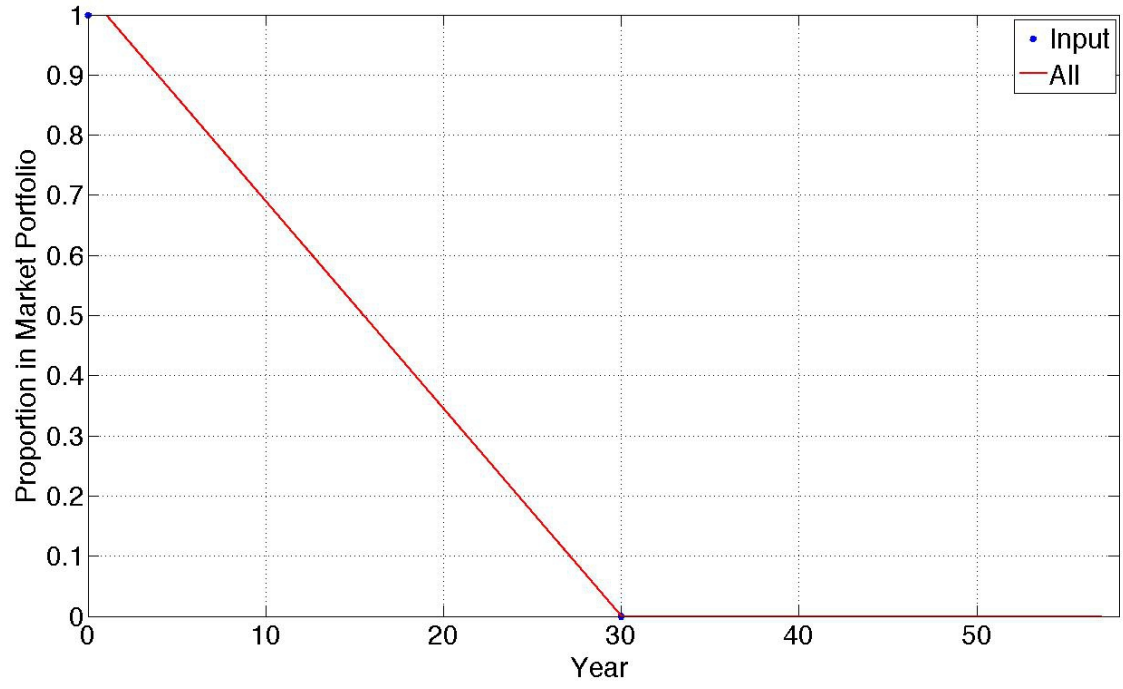
Thus far we have analyzed a case in which funds were invested in the market portfolio throughout the period covered. But our programs allow for the adoption of investment policies that change the mix of the market portfolio and TIPS from year to year. Practitioners typically advocate decreasing the risk of the investment portfolio over time, hence the term “glide path”, conjuring up notions of a plane gliding to ever lower altitudes. But if we wanted to do so, we could simulate a “reverse glide path” if desired, simulating a plane flying higher and higher.

To see the possible effect of the common strategy to decrease the proportion of the portfolio at risk over time, we set:

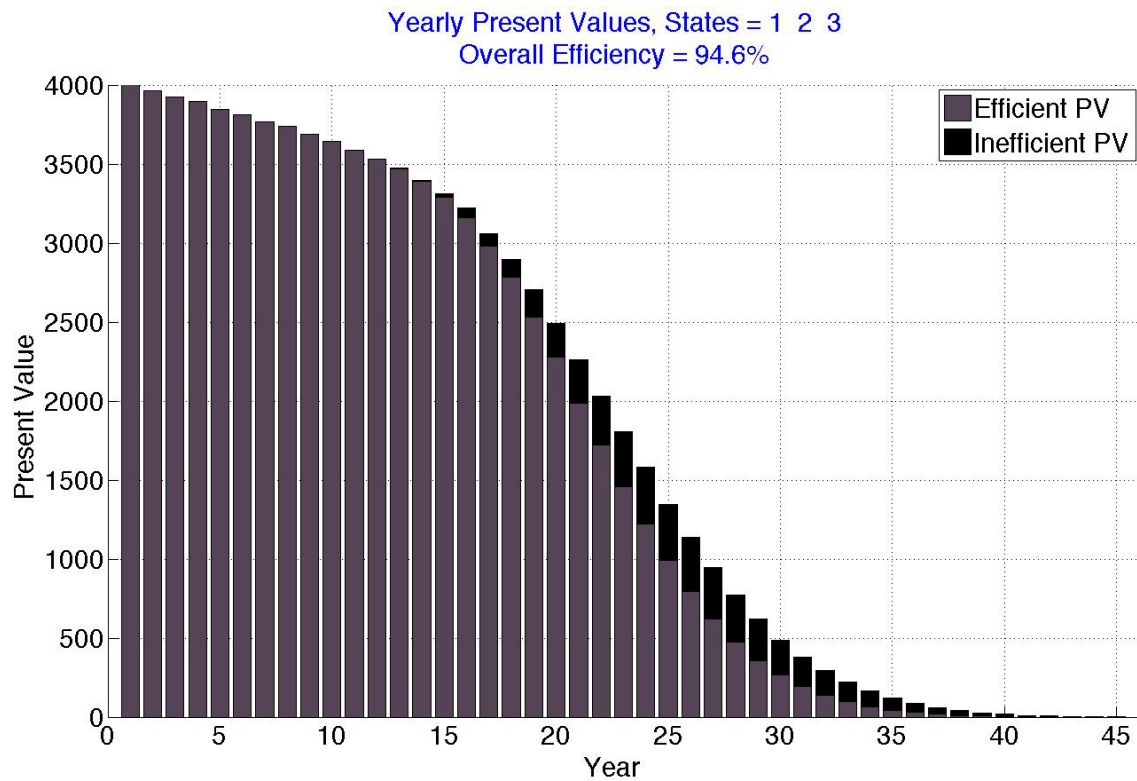
```
iConstSpending.glidePath = [ 1 0 ; 0 30 ];
```

producing a glide path that decreases portfolio risk until year 30, after which funds are invested solely in TIPS:

Glide Path: Market Proportions by Year



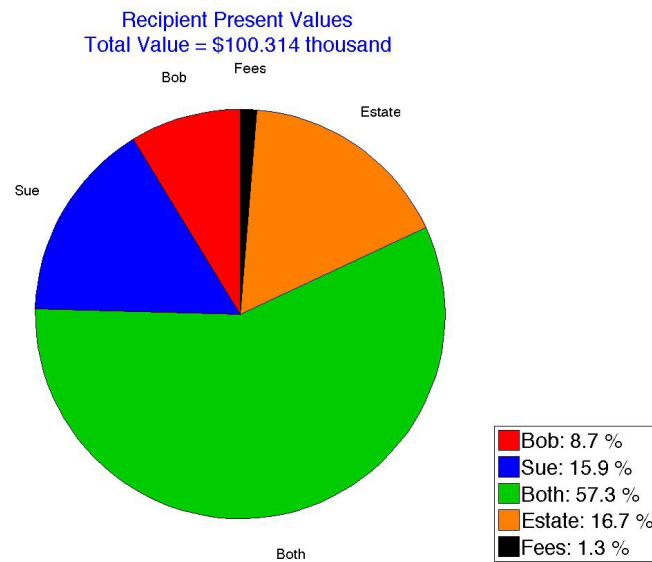
Whatever the desirability of other implications of the change, the impact on efficiency is definitely negative. Here is the relevant graph:



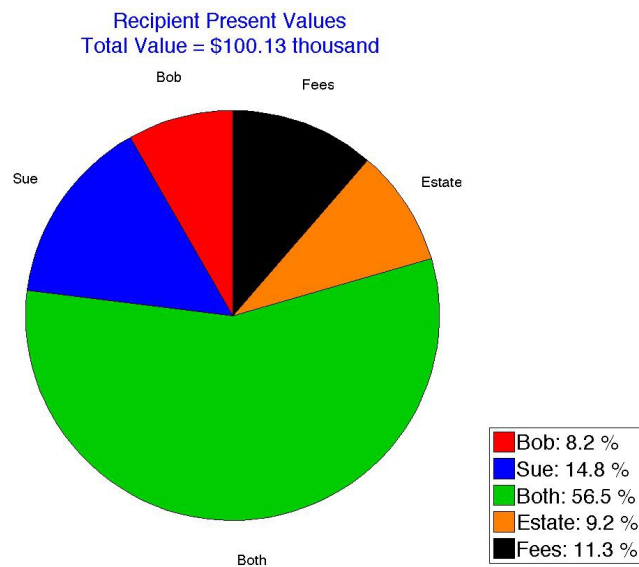
Note that the cost efficiency has dropped from 96.4% to 94.6%. We have made the results more path-dependent and are, in effect, getting even less for our money. As we will see in later chapters, there are better ways to provide retirement income.

Expenses

Thus far we have optimistically assumed that investment and any advisor expenses could be held to 0.1 percent of asset per year (10 basis points). But it is not at all unusual for advisors to charge 100 basis points or more per year and even use actively managed funds with additional expense ratios of up to 100 basis points. The impacts on retirement income can be dramatic. Here is the distribution of values with total expenses of 10 basis points per year:



And here it is with expenses of 100 basis points per year (with a slight difference in total value due to sampling error). The likely loss in income for poor Bob and Sue is equivalent to throwing away an additional 10% of their initial savings. Investor beware!



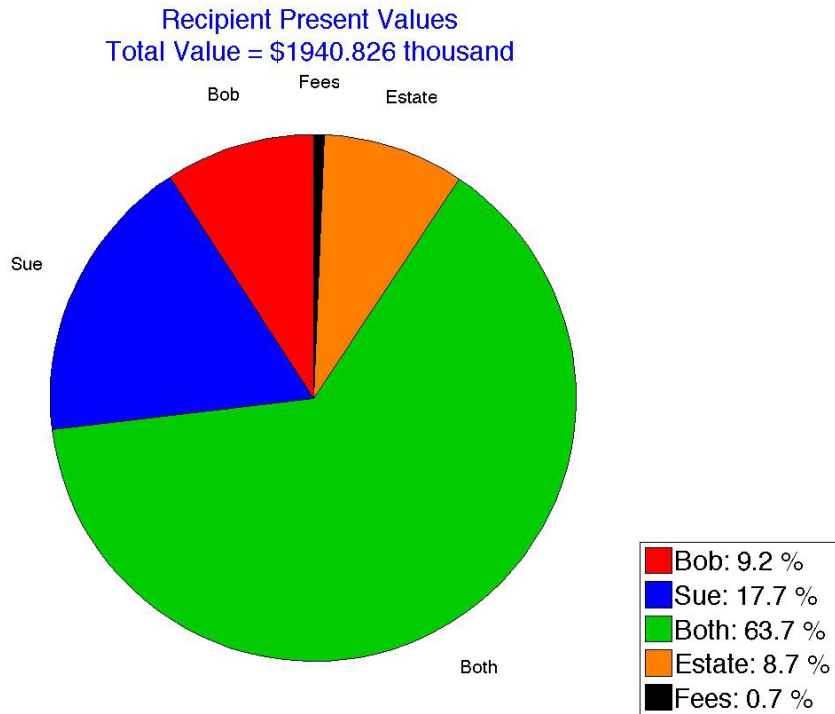
Constant Spending plus Social Security

Before moving on, it seems desirable to return briefly to the context in which most people evaluate spending procedures. For many, if not most, investors, discretionary savings are not the only source of retirement income. At the very least there is some sort of defined benefit plan, be it Social Security or a pension provided by a State or Local government employer. Combining such a plan with a constant spending policy provides outcomes that need not result in starvation if the latter “runs out of money”.

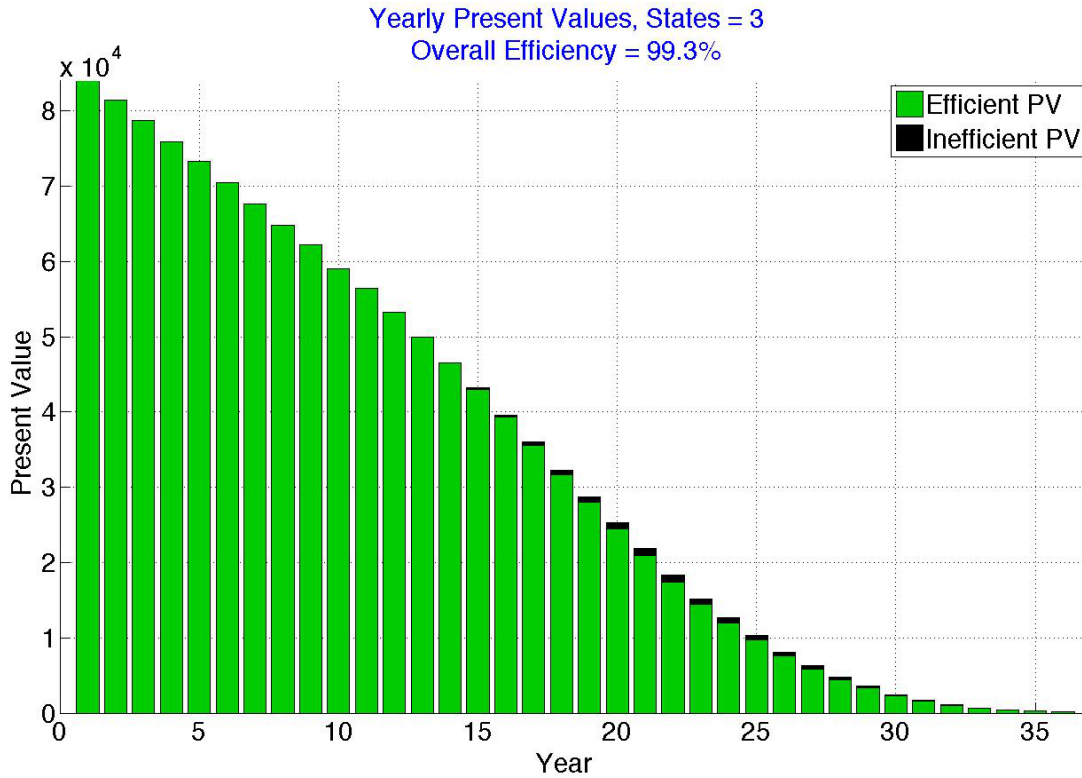
To illustrate, we return to Bob and Sue, assuming that they have the Social Security benefits we analyzed earlier plus \$1,000,000 invested in a constant spending strategy with a low expense ratio (0.10% per year) and yearly withdrawals equal to 4.0% of the initial value (unless and until the money runs out). Moreover, they have chosen to invest the latter entirely in the market portfolio in every year.

Here are the results

First, the distribution of present values. Note that the estate has a smaller proportion of the total, (8.7% instead of 16.7%) since it does not benefit from Social Security.

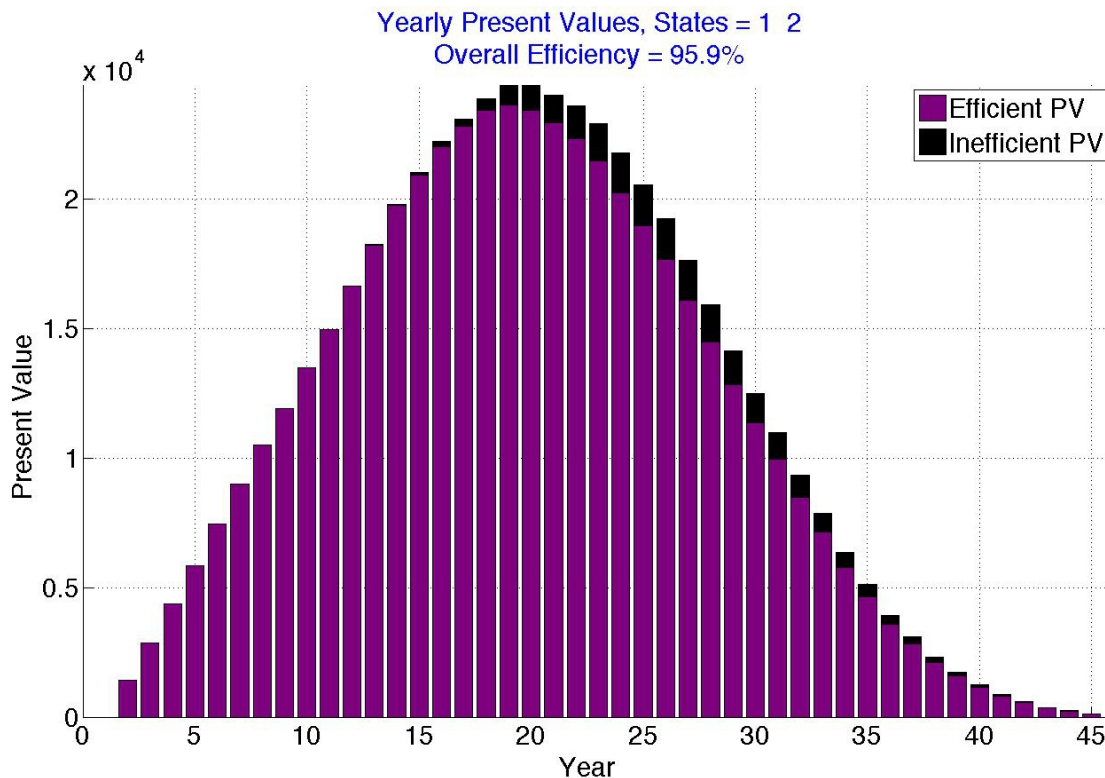


Next the efficiency. We now need to evaluate personal states 1 and 2 separately from state 3 since Social Security benefits differ when both are alive from those when only one receives benefits. Here are the results for state 3 (both alive):



Note that the efficiency is high since it is relatively rare for the constant spending strategy to run out of money when both are alive.

The picture is not as pretty for the states when only one person is alive:



Note, however, that the efficiency is nonetheless greater than it would be if the constant spending strategy were the only source of income, since Social Security is completely cost-efficient.

A video of the output from the case is available at:

http://www.stanford.edu/~wfsjarpe/RISMAT/SmithCase_Chapter17.mp4

The bottom line is that when a constant spending policy is evaluated in the context of other sources of income, it might not be as unattractive as when considered in isolation. That said, in the next chapters we will argue that there are better approaches to spending investment savings when annuities are not available, too expensive or subject to too much default risk.