AJO

WHAT INVESTMENT RISK REALLY IS, *ILLUSTRATED*

Strategic Asset Allocation and the Volatility of Assets and Spending During Retirement

Draft January 24, 2018

Investment advisors, managers, and their clients talk about balancing risk and return when advising clients about their strategic asset allocation (SAA) policy for the overall portfolio. This is true whether the client is an individual, a pension plan, a foundation or endowment, or a sovereign wealth fund.

SAA investment policy or strategy is typically set by choosing from among a series of different candidate asset allocation mixes located along some kind of efficient frontier, ranging from mixes with a relatively low expected return and a corresponding low level of risk, up through other mixes with higher expected returns and risks. So the advisor or manager discusses with the investor the tradeoff between expected return and risk, dialing them up or down together to a place thought to be consistent with the investor's tolerance for risk.

Yet, in our (not inconsiderable) combined experience observing our industry's practices, these discussions rarely focus on *what risk really is*. Often, even the investment managers or advisors themselves don't fully appreciate the nature of risk. Pension plans, for example, keep moving up the risk ladder, expecting to "get" the higher expected return that goes with it, and not understanding the very consequential meaning of taking on that additional risk, namely that you might get a much lower realized return rather than the higher, "expected" one. Individual investors are told something very similar, to hold lots of equities for their higher expected returns to support ambitious future spending plans, again without any meaningful communication about the higher risk that goes with that decision. What does it really mean "to take more risk"? Are these strategic asset allocation (SAA) decisions really made in such a way as to be consistent with the investor's risk tolerance, if the meaning of risk is not well understood? We'll focus here on the individual investor, but there are analogous discussions for investors of all types, of which we'll make brief mention where appropriate.

To be clear, our goal is not to discuss how to maximize the investor's utility when making this tradeoff; that can be left to other discussions. We're focused here on the task of simply understanding what risk really is, and how it changes with SAA policy choices and with different approaches to spending.

Along the way, we'll examine how risk looks for the traditional single-period capital asset pricing model, as well as for a multi-period consumption-based model, that is, a multi-period capital asset pricing model paired with a spending rule consistent with such a model. We'll also look at the impact of choosing a simpler spending rule, the so-called "4% rule;" it makes a difference to risk.

Asset-Only Risk, Under the Traditional CAPM

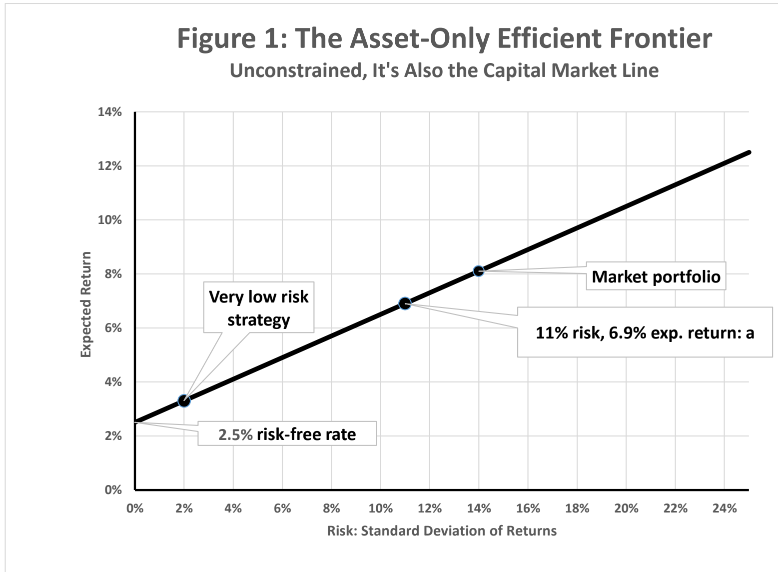

There are three levels to risk that we want to draw attention to. The first of them is risk in the context of the traditional capital asset pricing model (CAPM) of Sharpe [1964], Lintner [1965], Mossin [1966], and Treynor [1999], which seems to be the most commonly used SAA framework. In that model, we are accustomed to seeing a capital market line (CML) or an efficient frontier, representing the expected payoff for taking market risk. In the classic CAPM diagram in Figure 1, the y-axis shows expected return, and the x-axis shows expected risk, measured as the expected standard deviation of the distribution of future (realized) returns. As one moves from left to right along the efficient frontier (here, unconstrained, the efficient frontier is the same as the CML[1]), risk goes up, and the story — correct insofar as it goes — is that expected return also goes up as expected risk goes up.

The CAPM is an asset-only model, and it is pretty good, even with all its widely known shortcomings. But what does risk — frequently defined for convenience as expected standard deviation — mean to us, in this story? Usually, it is just glossed over, even in the context of this type of graph: "Oh yes, there is more risk as you move up and to the right on the CML, seeking higher expected returns. However, you *need* higher returns in order to have enough terminal wealth to meet your desired retirement 'number,' so let's move you out to the right to a higher return portfolio." No useful interpretation of "moving out to the right" is given, and rarely is any distinction made between *expected* returns and the risky realized returns that will actually be experienced. In other words, *the nature of risk isn't communicated at all*. Again, what does it really mean when we say that risk increases?

---

[1] We aren't showing the usual curved efficient frontier in Figure 1, which some readers will be looking for, and which would lie just beneath and tangent to the CML. Students of finance will remember that, with unlimited and costless borrowing and lending, the market portfolio can be levered up or down to create the linear efficient frontier shown in the figure; this is the famous Sharpe two-fund theorem, single period version. In effect, the efficient frontier moves up and becomes the CML. To simplify the discussion, we're idealizing a bit, using the two-fund version of the efficient frontier. As a result, the risk *vs.* return story that we tell here is a bit optimistic; while real world mileage will vary, the risk and return relationships will be directionally correct. Readers can readily adapt to their own discussion contexts.

What Risk Really Is, in Pictures

One way that we have had success in communicating the nature of this risk is to superimpose, over the CML graph, normal distribution curves showing graphically instead of just numerically the risk taken at various points on the CML, as in Figure 2.

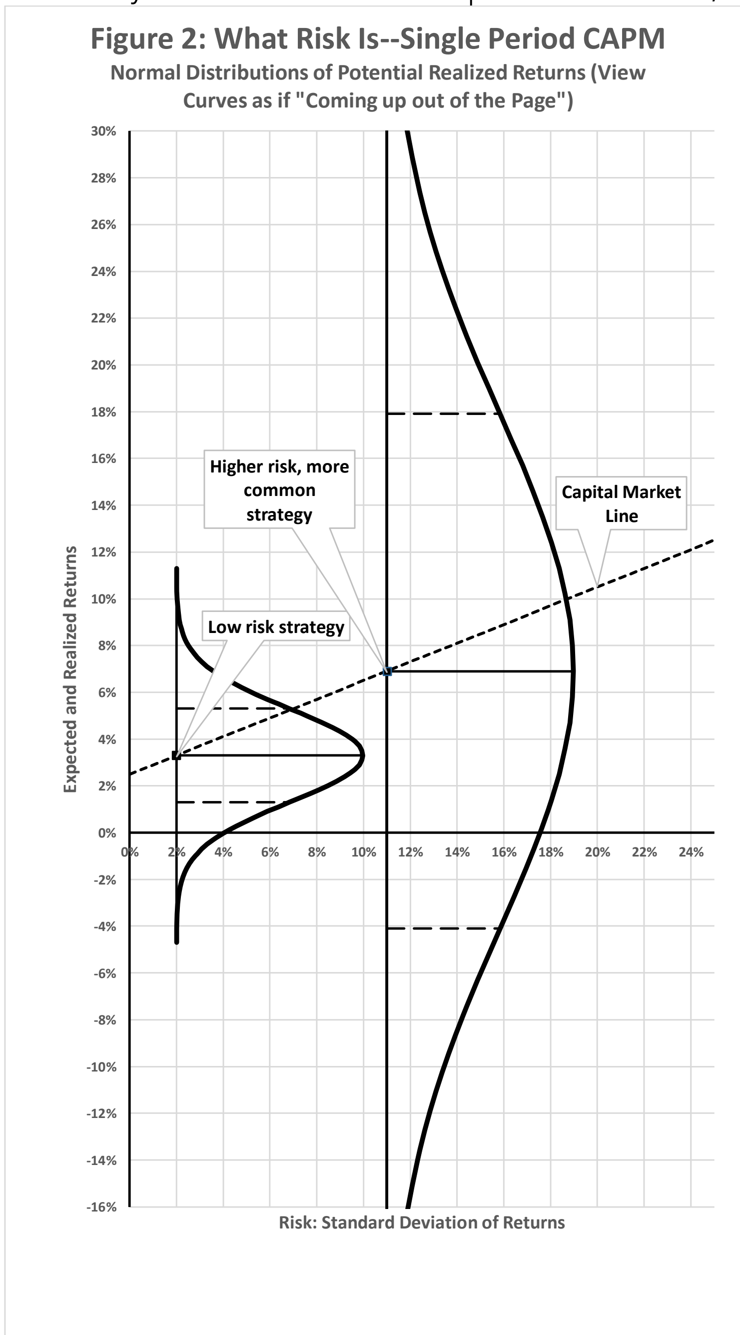

We have called out two different risk levels (of the many possible) to illustrate the nature of risk, one at a 2% standard deviation and one at a more commonly seen 11%, as shown on the risk (*x*) axis. Each of these has a normal distribution superimposed over the upward-sloping classic CML line, with a vertical size corresponding to its risk level. The base of each normal curve (the vertical straight line) intercepts the CML and the x-axis at the targeted risk level. Up and down, the normal curves are centered at their mean on the expected return or y-axis — at mean returns of 3.3% and 6.9%, respectively. The 11% risk SAA investment policy can be thought of as roughly similar to many of today's common 70% equity — 30% bond investment strategies. The 2% risk SAA investment policy would be considered very conservative.

You can think of these normal distribution "bell" curves as coming up out of, or perpendicular to, the page, making this a sort of *faux* three-dimensional graph: the normal curve's left-right dimensions have little to do with the dimensions of the left-right x-axis representing risk, so this notion that the curve is actually coming up out of the page into its own third dimension rather than spreading across the two-dimensional page is helpful in avoiding confusion about the meaning of its width.

The up-down or vertical dimensions (for returns) of the bell curves on the page do, however, correspond to the scaling of the y-axis. Recall that the x-axis is in units of standard deviation, representing risk; we translate the physical dimensions of those units to the y-axis in sizing the normal curve in order to show the range of possible realized returns implied by that standard deviation. Note that the vertical distance along the base for each of these curves from the mean to either the plus or minus one standard deviation line is exactly equal to the horizontal scale for risk — the 2% risk SAA investment policy has a 2% spread, or distance, from its mean to either its up or its down standard deviation marker; ditto for the 11% risk SAA investment policy. In other words, and importantly, the up and down dimensions of these bell curves on the page show the distribution of possible realized returns for that level of risk.

These first two figures have the same scale — but to show the distribution of possible outcomes, i.e., risk, we've had to show dramatically more of that scale for Figure 2. Figure 2 is so much taller than Figure 1 to accommodate the very large range of possible realizations. Risk is substantial!

The area between the normal curve and its base is the key to understanding probabilities. Specifically, the area between the base and the normal or bell curve is proportional to the probability of experiencing some particular level of realized return, or worse, for the area below that level (or, conversely, better than that level for the area above it). For example, note that the area between the base and the bell curve below the horizontal level of, say, the zero level of return — the x-axis itself — is much greater, as a proportion of the total area between the base and the bell curve, for the 11% risk level bell curve than for the 2% risk level bell curve. This higher proportion means that the probability of a negative return is much higher with the higher-risk strategy. In addition, the negative return is also likely to be worse (more negative) with the higher-risk strategy than with the lower-risk one. Look more generally — it is very plain that any higher risk strategy has many more disappointing outcomes than lower risk strategies, and many of those are more severely disappointing! In other words, when you take greater risk, the probability of doing worse during bad periods very definitely *increases*. And not only do you have a higher chance of doing worse, you might do *much* worse.

Let's state this as the important lesson to take home about SAA risk in this single-period CAPM: if you take more risk in order to increase your expected return, your realized return might be worse; in fact it might be much, much worse. But it also might be better — much, much better. On average, actually, the investor can expect it to be somewhat better; that's what it means to say the expected return is higher, and is the reason we might choose to take risk in the first place. But…there is that increased downside to consider — can the investor tolerate the pain, if "risk happens"?

This spread of possible realized returns *is what risk is*, at least for a single-period asset-only CAPM investor.

Let's remember how to use these two important words, so very useful when discussing what risk is: the *expected* return is just the mean or average of a very wide distribution of possible *realized* returns — and risk describes the standard deviation or spread of the *realized* return distribution. More risk means greater possibility of disappointment on the downside (and to be fair, greater possibility of reward on the upside). Moving up the curve "to get a higher expected return" exposes you to a greater chance of a favorable return, but also to a greater chance that you might do worse than you would have done had you taken less risk. We recommend to all that they adopt consistency in referring to expected returns and to realized returns using those terms; it goes miles especially when discussing risk.

Is the Single-period CAPM the Best Model?

This classic CAPM story of risk and return is extremely helpful in forming our intuition about the economics of investing, but it has some problems (actually many problems, but they are not that important to our larger story about the nature of risk). Not least among them is that the risk-free asset on the left-hand end of the CML is typically taken to be some form of riskless cash — usually short term government notes or bonds perhaps with one month or one year to maturity.[2] Yet we know no investor who moves his or her (or its) portfolio between cash and the stock market to make their asset allocation decision. Cash just isn't usually held in any significant amount (other than for market timing, not our topic today).

Also, the CAPM is what we call a "single-period" model. Usually we think of it as representing just the next period — say, the next year. But no real investor is thinking about just next year — they have a much longer horizon over which to invest, and more importantly to consume, their savings. Think about this limitation: It merely shows your expected wealth, in terms of expected investment returns, one year out. Even if we extend the asset-only model to multiple years, it still only shows the distribution of cumulative asset-only returns and wealth.[3]

But might this asset-only approach be missing something? What is the purpose of this investment fund anyway — just what is it supposed to provide for? Perhaps there is another deeper layer to the risk and return onion that addresses these concerns, that weighs risk and return in terms of one's intended use for the money.

---

[2] While Treasury bills are not issued with only one month to maturity, the Ibbotson studies (Ibbotson and Sinquefield [1976 *et seq.*] use existing bills with approximately one month remaining to maturity as their proxy for the riskless asset. This echoes many academic research papers working in this area.

[3] See Ibbotson and Sinquefield [1982], for their approach to showing the distributions of future returns and portfolio values over multiple periods in an asset-only forecast based on compounding up single-period expectations.

Multi-Period Capital Asset Pricing Models

In the early 1970s, Robert Merton started writing about a new and much more sophisticated approach to modeling capital asset prices, using "multi-period stochastic dynamic programming" mathematical technologies. Anything with such an unintelligible description is in fact likely to be inaccessible to most mortals, so perhaps it is no surprise that study and refinement of his approach has been largely confined to academic researchers rather than becoming part of the practitioner's and investor's worlds. His approach is in fact quite complex conceptually — but it is also a vastly improved way of thinking about how investors should set their SAA policy to maximize the value (expected utility) of their potential future consumption over many future periods.

This approach has virtually taken over the advanced academic study of investment strategy and the market pricing of assets. Literally hundreds of papers, many quite important to our points here, have been written exploring a multitude of permutations of his work, all using a version of multi-period stochastic dynamic programming to frame the decision. And all of them are also largely unreadable except by other academic economists and a handful of very sophisticated investment practitioners.[4] We'll cite only Merton [1973] as the best example of the origination of this line of thinking, and Merton [1990] as the exemplar for the consumption version; the others contributing important pieces to its development are too many to cite here.[5]

We'll refer to all of these generically as "multi-period capital asset pricing models", or "multi-period CAPM".

What Does a Multi-Period CAPM Spending Plan Look Like?

The big insights from Merton and the remaining literature that follows, if we can be allowed some license in paraphrasing, are (1) that assets are held for a reason, to fund some plan for spending or consumption in the future, and (2) that *risk* is measured with respect to the certainty or lack of certainty of there being funds available to support whatever future spending the assets are being held to provide. And, furthermore, this future spending plan is likely to be a multi-period rather than a single-period plan, so it needs a multi-period approach to choosing one's risk *vs.* return trade-off.

There are a few things that come out differently in the multi-period SAA world than in the single-period world. First off, we treat the multi-period liability-hedging or consumption-hedging asset, instead of "cash," as the risk-free asset, in service of the idea that riskless, or minimal-risk, investing really means protecting consumption or other spending over a multi-period horizon. Think of this rate as the average expected return on a ladder of cash-flow-matched bonds whose coupons and principal payments would, if they constituted 100% of the assets, precisely fund the spending plan and

---

[4] For those more mathematically comfortable practitioners desiring to get a better handle on this technology, sometimes referred to as the "lifecycle model," we strongly recommend the excellent survey text by Charupat, Huang, and Milevsky [2012]. It is reasonably accessible — and meaningful. Also, the CFA Institute Research Foundation has published a readable summary of this thinking; Ibbotson, *et al*, [2007].

[5] Waring and Whitney [2009] replicates many of the useful characteristics of these complex models in a comparatively simple mean-variance framework, capturing the multi-period consumption optimization problem in a fully tractable and quite informative way, and clearly demonstrates the nature of the two-fund theorem in a multi-period context.

because of the match, thus hedging it against risk from interest rate changes. Why does this hedging portfolio make spending risk go away? If rates go up, driving the portfolio's value down, the higher rate still provides the same planned spending from the new lower-valued portfolio (and vice versa) because of the duration and/or cash flow match.

To keep the SAA portion of the discussion simple, we'll think of it in a multi-period version of the two-fund theorem — the investor is choosing the relative proportions of (1) this multi-period consumption hedging risk-free asset (rather than cash), and (2) a second portfolio of risky assets intended to proxy the CAPM's "market portfolio." Even when not constituting the entire portfolio, the hedging characteristics of the risk-free asset in this two-fund approach will reduce consumption volatility relative to a SAA policy developed in a single-period CAPM framework — an immediate benefit of moving from single-period to multi-period space.

Maximizing the expected utility of consumption in a multi-period SAA problem implies some specific characteristics of the spending plan. It implies that the assets shouldn't run out too soon or too late: This in turn implies that at every point in time the present value of planned future spending must be equal to the value of the assets at that point. Economically, the liability can never be greater or smaller than the assets, as noted by Waring and Whitney [2009]. Think about it: you can't spend more than you have, although with high *expected* return investments and some good luck you can increase your *expected* spend over time. And you will spend everything you do have, even if it is in the form of unanticipated gift to heirs (or even escheat!); you may as well try to maximize the utility of that spend by focusing it on your priority spending plans and making an intentional contingency plan — a bequest — for the fair likelihood that you won't live for your full maximum retirement horizon leaving some portion of the portfolio unspent on consumption and available to heirs.

Keeping the present value of future spending equal to assets is not inflexible. It leaves room to "shape" the spending plan, which can be a level payout in *nominal* terms over the planning horizon (shrinking in real terms); it can be level in *real* terms (growing with inflation, thus providing level spending power); it can be "bent" to spend more in the earlier years of retirement and thus less in the later years; or vice versa — less earlier and more later. There is an infinite number of possible shapes for the spending, as long as the present value of the spending matches the value of the assets at all times. (These reshapings are accomplished in the multi-period CAPM literature in different forms, but with generally similar effect, using parameters known as the "subjective discount rate" or the "personal patience rate" or the "intertemporal elasticity of substitution"). Reshaping spending and consumption is an interesting topic and we'll publish more on it, but we have previously touched on it in our spending rule article, Waring and Siegel [2015]; Milevsky and Huang [2011] have done interesting work studying the effect of longevity risk on the reshaping of spending over time.

Visualizing Spending Risk in a Multi-Period CAPM

With that introduction, let's examine the second level of risk of interest to us today, and examine how investors can visualize the nature of the risks they will face as they consider alternative SAA policies when considering how best to maximize the expected utility of their multi-period consumption. Maximizing utility means something like this: Maximize the expected return of the assets for the purpose of increasing expected consumption spending, at *no more than an acceptable level of unpredictability of that spending from period to period — an acceptable level of spending risk*.

As we said above, assets are always held for a reason and investment policy should be decided relative to that reason. For individuals, the biggest and most important reason might be for retirement income. But it might also be to provide a contingent reserve for emergency medical or other expenses, to fund a child's education, to provide funds to a favorite charity, foundation, or school, or to provide a generous inheritance to children or friends; one can imagine many other investment goals that one might want to provide funds for.[6]

For the assets of an endowment, the *raison d'être* is to provide operating support for the institution and perhaps scholarship money; similar for a foundation. For the assets of a pension plan, it is to pay benefits to retirees. For most of these types of investors, we can identify the risk-free liability-hedging asset, and thus see a way to hedge most or all of the market risk out of the particular spending plan.

But we promised in this article to focus on a particular goal, the retirement income needs of the individual investor; this keeps it simple because an individual needs a single string of cash flows for spending, but the nature of risk would be similar for a different or more complex set of goals.

So let's assume that our investor is retiring today at age 60, and has a rather simple, single spending goal: She wants to plan to spend her savings of $2,525,427 in a manner that gives her a payment each year, payable at the beginning of the year, over what she feels is her maximum horizon of 45 years, living to a possible maximum age of 105; she doesn't want to risk running out of money, preferring the risk of dying without having spent it all yet and letting the otherwise unspent assets serve her bequest goals, which she feels are strongly secondary to her own support. She has choices for different spending shapes, as discussed above, but she chooses a level nominal spending pattern. In real, inflation-protected terms, this means spending somewhat more, earlier, and somewhat less, later.

The spending plan that meets this goal and is also equal in present value terms to her portfolio is thus a 45-year level nominal annuity payout stream. Not a commercial annuity, but a self-paid "virtual" annuity, no commercial annuity provider being involved.[7] So ordinary annuity mathematics can tell us what she can afford to spend each period.

---

[6] These goals are all contemplated in today's goals-based SAA investment policy models, but usually asset allocation policy for each goal is determined separately one bucket at a time; we believe that asset allocation policy is more optimally determined in an all-at-once full surplus optimization along the lines suggested in Waring-Whitney 2009.

[7] Because of high fees and risk charges and life tables biased to protect against selection bias, the possibility of annuity provider default due to typical insurer asset-liability mismatch over the very long future time periods involved, the inflexibility of SAA investment policy (implicitly an annuity is an all-bond portfolio, and no, a variable annuity is not the answer to this), and lack of liquidity (you can't cancel the annuity and get your money back without substantial costs), our investor does not desire to purchase a commercial annuity. We're not sure that she's wrong; annuities are good, even excellent *in theory*, but today's commercial annuities are far from ideal.

The yield on our hypothetical multi-period liability-hedging bond is assumed to be 3%, a modest half-point more than the one-year riskless rate we used in the single-period examples of Figures 1 and 2, and naturally it is a different risk-free asset than that used in those earlier figures, a multi-period consumption-hedging ladder of cash-flow matching bonds. Our investor is considering a risky SAA investment policy consisting mostly of risky assets with a smaller allocation to the risk-free asset, having an expected return of 6.9% per year with an 11% standard deviation — similar to the risky portfolio highlighted in Figures 1 and 2.[8]

What is the risk exposure to her planned consumption that she faces with this somewhat risky but typical SAA investment policy?

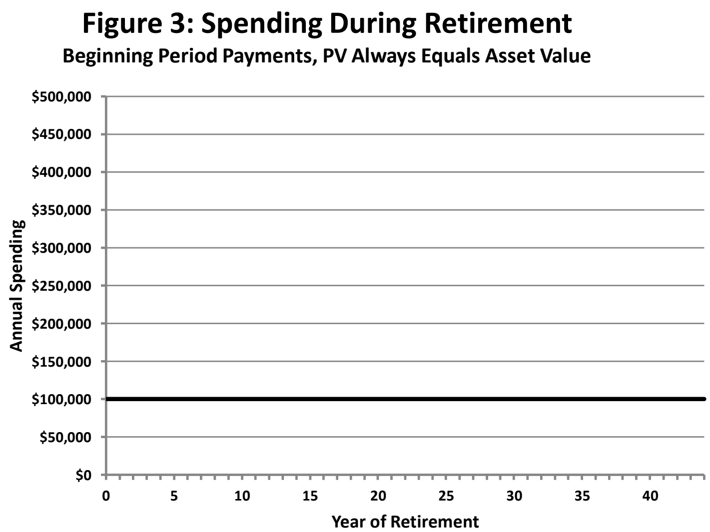

First, let's review that if she were to invest all of her money in the multi-period riskless or liability-hedging portfolio, a ladder of bonds with interest and principal payments that are cash-flow matched to each of her next 45 annual spending payments, she would have virtually no risk to her spending plan.

The annuity payment calculation uses the current interest rate on the risk-free asset, the portfolio value available (PV of spending must equal asset values), and her retirement time horizon.

Putting this into Excel™ annuity payment notation to summarize the problem:
Annual Spending = PMT(3%,45,2525427,,1) = $100,000. She could spend $100,000 per year, safely, every year as shown in Figure 3.

Her portfolio would decrease slowly over time as per Figure 4, not running out of money until the planned time.

[8] In this article, returns and standard deviations are presented in discrete, not continuously compounded form, in order to be comparable to asset allocation practice; for our simulations they are appropriately converted.

Figure 4: Asset Value Over Time, Level Spend

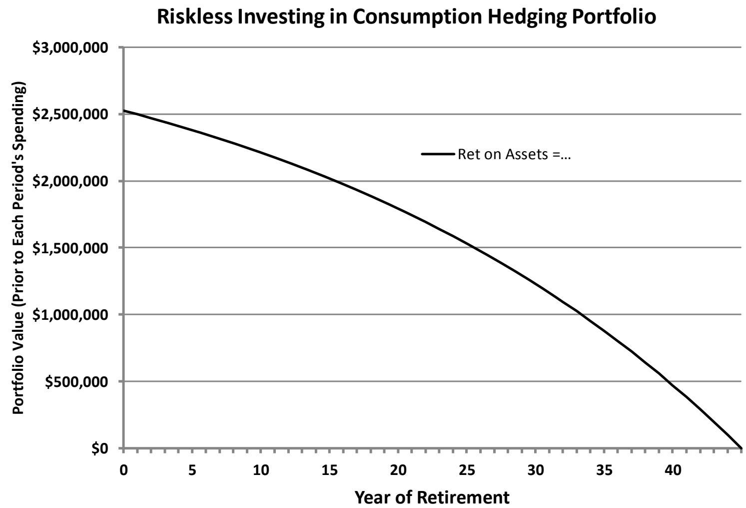

Of course, government bonds with such nicely matching maturities are not available, so our perfect liability hedge will in reality be somewhat less than perfect. For example, instead of a cash-flow-matched portfolio, she could invest in a portfolio that combines duration-matching as well as elements of cash-flow matching as a second-best alternative, and it would do an imperfect but pretty darn respectable job of hedging her spending, keeping the risk to spending quite low.

What About an Asset Allocation Policy That Includes Risky Assets?

Ok, we have established spending under the multi-period risk-free consumption hedging benchmark. But she is considering a modestly risky SAA policy. Her motive in considering this strategy would be the hope that the higher return expectations would translate into a larger portfolio and in turn greater expected spending than under the risk-free benchmark. What is the risk that such a higher-risk policy would introduce to her spending, relative to a fully riskless hedged position such as in Figures 3 and 4, or for that matter relative to any other SAA policy?

We can't demonstrate this multi-period risk with single normal curves, as we did for the other, less complex, single-period example shown in Figures 1 and 2. We can, however, use Monte Carlo simulation to understand both the distribution of returns of the risky portfolio and the distribution of resulting spending plans over time, to get similarly useful graphical descriptions of risk. In our simulation we will generate 12,000 "runs" of realized returns of our portfolio in each of the 45 periods, creating 12,000 "runs" of random returns, each run a plausible "future realized return history". We assume that each year's return is independent and random, thus producing a lognormal distribution of returns.

From returns, we create portfolio valuations: From the 12,000 random returns generated in the first period, we apply them to the initial portfolio value to generate a new set of 12,000 period-end portfolio values for use in starting the next period. From each of these portfolio values in turn, we calculate the 12,000 appropriate spending payments for that year using annuity math; more in a moment. We'll subtract that payment from the assets to get the 12,000 portfolio values at the start of the next period. Continue the process, in the following year each of the portfolio values will be subject to 12,000 new random returns, that year's appropriate spending, and a new period-starting portfolio value; repeat for 45 payment cycles. From these 12,000 portfolio valuation and spending paths over each of the 45 years, we pull a series of percentile values and plot them against time.

In a nutshell, the spending calculation is Waring and Siegel's [2015] "annually recalculated virtual annuity" (ARVA) method of setting the spending rate each period in a portfolio with a fluctuating random value. It's similar to the annuity method used above for Figures 3 and 4, but because the portfolio values fluctuate, we have to recalculate the payment every period using the then available portfolio value, and a spending time horizon reduced by one period. This constantly ties the present value of the future spending plan to the actual but volatile value of the portfolio. Using ARVA to calculate the spend, one never runs out of money and never leaves money unspent (if one lives for the full consumption planning horizon).

Absent portfolio fluctuations, the spend would be level in nominal terms, as specified by our investor (it could have taken many other shapes, at the wish of the investor, as discussed above); if the portfolio fluctuates, the spend must fluctuate. This is because, all else being equal, an annuity payment is always directly proportional to the principal value.

Figure 5 shows the distribution, by percentile, of the value of the portfolio over the 45-year time horizon. We see that there is a large disparity between the higher and the lower percentiles—from extremely poor wealth outcomes to extremely wonderful wealth outcomes—depending on the runs of good or bad luck experienced by the portfolio in the markets in the 12,000 different simulated realized-return histories. Because of the way we calculated the spend, the portfolio of course goes to a zero value in all cases by the end of our horizon.

Figure 5: Portfolio Value Over Time, Level Spend

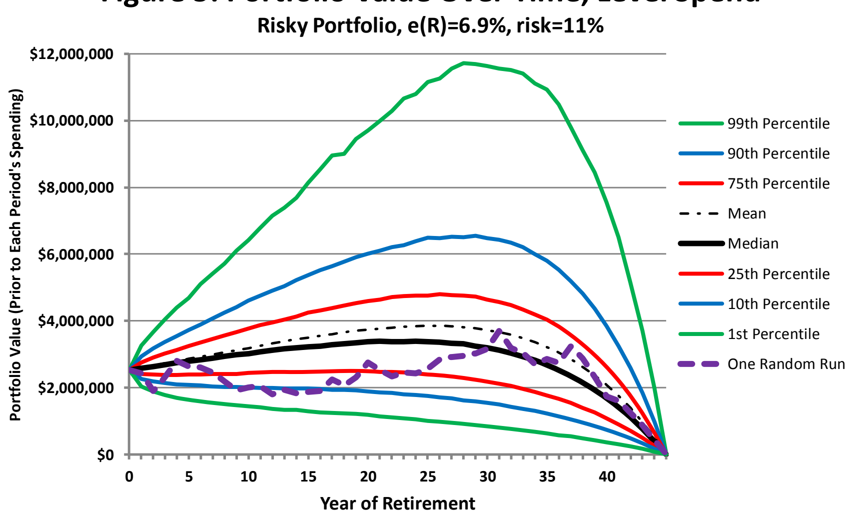

Note that we have also drawn in a line representing "one random run," just one of those 12,000 randomly generated future histories in the Monte Carlo simulation data. This should be helpful to those new to percentile distribution charts, which are in some percentile sense an average of all of the different runs generated in the Monte Carlo process. Some of these runs ride more or less in the middle of the distribution, but many more have wild swings or show long runs of good or bad returns. In this particular run, one can intuit, from the downward-sloping trend towards lower percentile positions during the first 10 years that the portfolio encountered a period of returns below expectations; it then levels off with realized returns closer to expectations and holds its percentile value for a while. Later, maybe after year 18, you can see that it enters a period of better realized returns where the portfolio grows in value relative to the percentile lines, returning to levels at and just above the median and mean levels as it levels off again in the final years. We chose to show this particular run because it allows us to illustrate how uncertain and volatile returns flow into the portfolio and thence into the spending levels — which takes us to Figure 6.

Figure 6 shows what happens to spending as a result of the gyrations of our investor's portfolio. Not surprisingly, in the wonderful wealth outcomes represented by the upper percentile asset value ranges, spending can safely go very high. But in the poor wealth outcomes, spending must stay low in order to safely last through the 45-year horizon.

Figure 6: Spending During Retirement

Beginning Period Payments Given Risky Portfolio

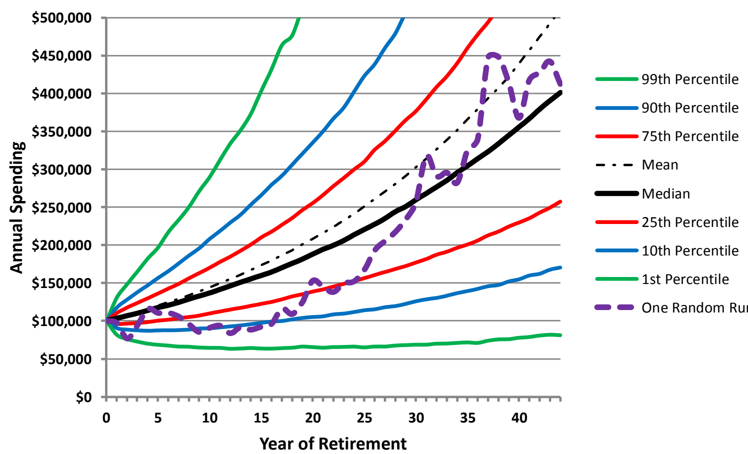

How comfortable would the investor be with the bet on this riskier portfolio if she were to choose it? Let's consider her spending level and volatility: Spending is dominantly (the great middle and upper part of the percentile distributions) well above the $100,000 per year benchmark that our investor would have safely gotten if her SAA investment policy had been to invest in the riskless portfolio of hedging bonds. In fact, we had to seriously truncate the top of the chart, where spending is many, many multiples of the benchmark. Even at the median, spending is growing rapidly throughout the time period. Wonderful. So far.

But there is a significant chance of not being able to spend even as much as the $100,000 per year benchmark — during extended periods of disappointing market realizations she would have to reduce her spending below, even far below, the $100,000 level that she could have achieved by holding only the riskless asset. During the first five years, there is a 25th percentile probability of a small disappointment. At the 10th percentile probability she might be spending significantly, but not penuriously, below the benchmark during the first 15 or so years. But between the 10th and the 1st percentiles she might be spending quite a bit below benchmark, possibly for the entire retirement period. And it is even possible to underperform the 1st percentile, experiencing realizations that put spending well below the lowest line. So there is some real risk of reduced, rather than increased spending.

Again, in addition to the percentiles, Figure 6 shows spending for the same single random run that was used to generate the one random line for the portfolio value in Figure 5. This connects the dots for how the volatility in the portfolio value shown in Figure 5 flows through to spending. Improving portfolio values lead to improving spends, and vice versa. One can imagine that the spending level would have to decline during the poor returns of the early years, and it does. And of course, as the portfolio recovers, spending improves, a constant process of adjustment of spending to match the portfolio value. Spending recovers to somewhat above the median in the latter years. It all ties out perfectly. In each of the other of the 12,000 realized return histories, spending will increase or decrease as these volatile realized returns are experienced by the portfolio.

Informing the Strategic Asset Allocation Decision

Back to our question: Is our investor willing to take this risk? Does the value of the possible wonderfully high spending outcomes outweigh the pain of the possible more penurious outcomes? Would she prefer the safety of the riskless hedging portfolio? Or would she rather be somewhere in between, maybe even with the tighter but generally lower spending distributions that would be associated with a very low 2% level of expected risk in a portfolio with more hedging assets and fewer risky assets (similar to the alternate portfolio in Figure 2)? What level of risk maximizes the expected utility of her future spending during retirement? Examining a graphical description of spending risk might help inform an investor's answer.

Let's look at the spending associated with that lower risk 2% standard deviation portfolio. To economize on space here, we'll skip the portfolio value chart comparable to Figures 4 and 5, but it suffices to say that the portfolio value distributions are much, much tighter than Figure 5, and trend much more like Figure 4; no surprise, given the low risk.

Figure 7 lets us review the volatility of spending under this much more conservative asset allocation policy. Likewise, it is a much tighter distribution. The investor is taking some amount of risk in the hope of some small but appropriate level of expected return, so it is no surprise that mean and median spending levels do increase above our riskless baseline of $100,000 per year. But it won't be a surprise that the improvement is not nearly as much as in the higher risk, higher expected return case. Notably, the downside risk distribution is also quite tight, spending never dropping below $83,000 at the 1st percentile level — the comparable low spending value in Figure 6 is much lower, at just over $63,000. There has been some very significant downside risk reduction.

Figure 7: Spending During Retirement

Very Low Risk Strategy, e(R)=3.3%, risk 2%

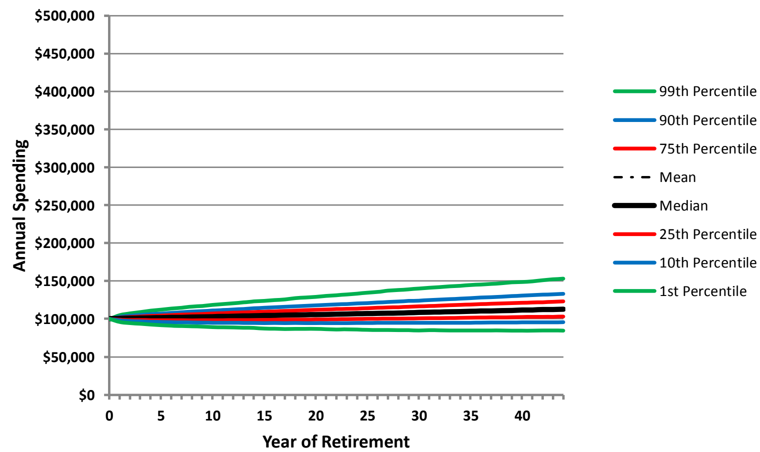

Between Figures 3, 6, and 7 we see how spending risk changes with more and less aggressive SAA policies, giving wider or narrower distributions of possible spending outcomes. For the multi-period investor funding a spending plan designed to last the full retirement horizon, these distributions of spending outcomes over time are what risk *is*: Greater SAA investment policy risk means greater spending volatility over time, including greater risk of spending below the risk-free benchmark spending level.

If the investor is considering the 11% risk SAA policy of Figure 6, but the spending volatility at that risk level is uncomfortable, the answer is to adopt a more conservative, more fully hedged SAA investment policy (fewer risky assets, more risk-free liability-hedging assets), which will tighten up the percentile ranges. It probably won't be as low risk as our 2% risk portfolio shown in Figure 7, or the risk-free portfolio in Figure 3, but somewhere in between the extremes.

Risk Introduced by an Overly Rigid Spending Plan: The 4% Rule

Let's move on to the third level of risk that we want to explore today. It's a little different, being more about how risk increases with a poorly designed spending plan, and less about the relationship of risk to the SAA investment policy. Yet our tools for understanding risk caused by investment policy in a multi-period retirement framework are perfectly suited to the task of analyzing risk created by a spending rule.

When reviewing the volatile spending levels that go with the volatile portfolios (Figure 6, above), one wouldn't be blamed for asking whether or not this volatility could be "smoothed away" in order to make the investor more comfortable. We answer that question by examining how the risks change when using the biggest smoothing rule of all, what came quickly to be called "the 4% rule." We're going to conclude that the faults of the 4% rule as a spending rule always add risk relative to a multi-period CAPM-inspired ARVA rule.

The word on the street is that the 4% rule is used by many or maybe even most financial planners and advisors in helping their retired clients plan their consumption. This spending rule, in its original form, says to spend 4% to 5% of your beginning portfolio in the first year of retirement, and to increase that first year's dollar amount by inflation for each subsequent year's spend. Its originator, William Bengen [1994], thought, after carefully and systematically reviewing the actual return histories of the U.S. markets (using Ibbotson data), that this rule was "safe" for a retiree with a 30-year (minimum) spending horizon.

In effect, this is a permanent smoothing rule, an effort to find one spending rate for all time despite holding a random walk volatile portfolio.

Bengen didn't stop his careful efforts with the original 4% rule, but continued to develop his thoughts, coming to endorse 4.5%. Recently, he acknowledged some of the problems and recommended that if withdrawals became too large (small) a percentage of the remaining assets due to poor (good) market conditions, it would be wise to reduce (increase) spending: While he had originally thought a reduction in spending was "anathema," he had come around to realize it could be necessary in worst case (best case) situations [2012]. Even in its updated version, the 4% rule is still fairly characterized as a smoothing rule, but with occasional adjustments.

For the moment, however, we'll stay with his original, strict 4% rule rather than his more developed approaches, as it is the one that is so very widely used: It is one of those catchy, simple, and easy-to-remember rules that many investors and advisors have seized upon as if it were the full story and nothing but the full story, even ignoring Bengen's ongoing efforts to improve the rule. We'll explore it with our risk graphics.

But, first, for mental calibration let's note some differences between Bengen's 1994 work and the way we have approached the problem of setting spending, using ARVA. He appears to have been motivated to find a level, smooth, and unchanging (other than for inflation) spending rule, which limited his choices of approach. He used a less aggressive portfolio than we are using, some 53% in equities and other risky assets versus something in the range of 70% for our work, which we believe to be more representative of actual investor behavior today (he may well be correct that the less aggressive asset allocation is more appropriate for many or most retirement investors).

Further, the 30-year retirement minimum horizon of his study seems to us too short to be safe, so our spending horizon is longer than his was: The Social Security Administration's life expectancy tables now go out to age 120, showing a vanishingly small but not impossible length of life; at age 105 for men and 110 for women, the probabilities are small but not vanishingly so. A retiree at age 60 or even age 70 might well outlive Bengen's 30-year horizon; thus we used 45 years in our work here. It seems to us prudent to plan the longest possible life horizon, biasing toward the risk of our money outliving us rather than of us outliving our money.[9] This shorter planning horizon

---

[9] In other work, one of us (Siegel), with co-author Stephen Sexauer, has recommended using deferred life annuities to hedge the tail risk of possibly living a very long time. See Sexauer and Siegel [2013]. The other of us (Waring) consistently argues that despite theoretical advantages the value of such annuities in the real world is overstated for the reasons given in footnote 7, sufficient to make ARVA self-annuitization the more attractive alternative. (Vive la différence!)

baked into the original version of the 4% rule has been completely overlooked by the advisory community — they treat the rule as if it were for the maximum possible life without caution for the possibility of one's life happily lasting much longer than 30 years; you'll only occasionally see discussion about changing the rule for investors who might have a longer or shorter horizon.

Lastly, Bengen made the conscious choice to do a "deterministic" study of spending outcomes based only on the one run of historical returns available to us, rather than a "stochastic" study using a good model of future returns, such as the widely-used random walk model we have used here for our Monte Carlo simulations. Thus he didn't have a chance to look at the entire distribution of investment return and spending possibilities; his options were limited by the choice of method.

To compare the 4% rule to our results above using an ARVA spending rule, we ran it through the same asset return simulation that we used for the ARVA examples above, increasing spending each year by an assumed constant 2% inflation rate.

In this case, for the same starting portfolio as the one we looked at earlier, Figure 8 shows that the initial-period 4% spend would be $101,017, growing after that at our assumed constant 2% inflation rate to $241,436 by the time the 45th annual spending payment is taken out (at the beginning of year 44). Note that the 4% rule gives in this case very close to the same first year spending that we planned for using our ARVA spending rule method, $100,000. This is largely coincidental, a confluence of the interactions of time horizon and our assumed discount and inflation rates; with different assumptions, the two could be substantially different.[10] However, as shown in Figure 6, at the median for the 45th year our plan would be spending just over $400,000 and at the mean over $500,000, far greater amounts than the $241,436 maximum anticipated at that time under the 4% rule. Despite starting at the same level, in the mid-grounds of probability the 4% rule produces much less spending than the ARVA approach over time.

---

[10] Higher or lower discount rates are very important to the baseline spend in our ARVA spending rules, based as they are on annuity mathematics, and over a quite long horizon so that the interest rate makes a substantial difference. But there also is another notable assumption difference: the 4% rule bakes in planned inflation, here assumed to be 2% per year, and our examples do not (they could have). In our examples, we assumed a constant *nominal* spend as shown in both Figure 3 and Figure 6, this constant spend altered in Figure 6 only by the variability of the portfolio's value.

We can tell you how this would have worked out for Figure 6, however, if instead we had biased our spending to grow with inflation. Planning to grow spending with inflation implies spending less earlier, and more later, in order to preserve the equality of the present value of spending with the value of the assets at all times. Our Figures 3 and 6 both would show consumption beginning at $69,001 in the first period, much lower than under the 4% rule. In the riskless case in Figure 3, the spending path would remain well below the planned level shown here in Figure 8 for the 4% spending rule — but it would be truly safe from complete spending ruin, where the 4% rule has a very significant chance of spending ruin. On the other hand, the path for the median in the risky asset case shown in Figure 6 would end at $661,972, far above the 4% rule's expected ending spend in Figure 8, and well above the ~$400,000 ending spend in Figure 6. Again, the overall advantage goes to the ARVA spending rule rather than the 4% rule.

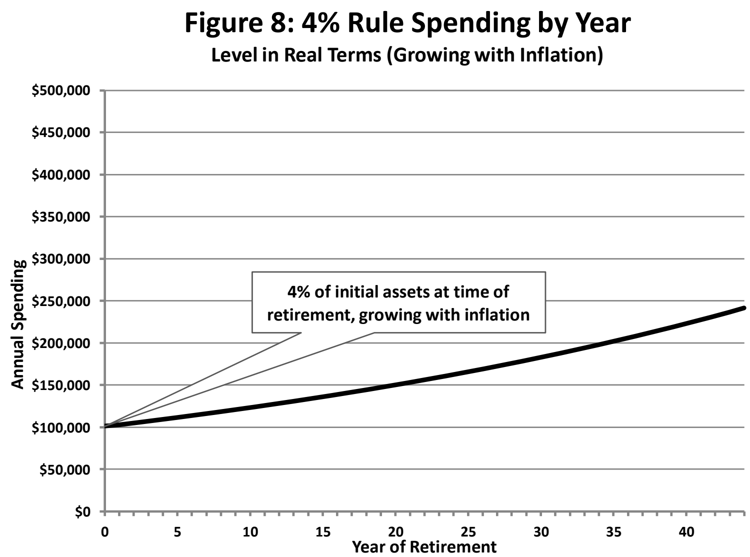

Figure 9 shows the problem graphically in a picture for 4% rule enthusiasts to study carefully: the steady and smooth spending rule shown in Figure 8 results in wildly diverging portfolio values. When markets have a long period of excellent returns, the portfolio's value goes through the roof, with no corresponding increase in spending, and when there are long periods of poor returns, the money runs out long before the end of our planning horizon — even at something greater than the 25th percentile of probability. Of course, when the money runs out, the nice smooth spending line of Figure 8 stops abruptly, dropping to zero. This can't be seen as acceptable — running out of money is not a very attractive option. One sees why Bengen feels today that it is necessary to reduce (increase) spending when things go poorly (well) for a while; he's studied this extensively, and we agree with him although we would review and reset each year in an ARVA framework in order to tie the spending plan to the assets so that they are equal in present value terms at all times.

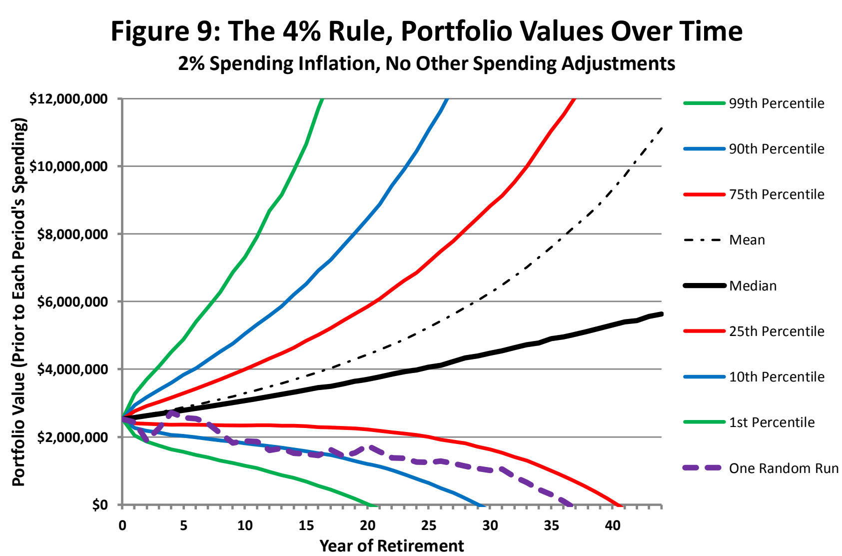

We may worry more about running out of money than leaving too much to our heirs, but if one contemplates Figure 9 for more than a moment one is struck by the fact that the outcomes are biased very strongly towards leaving a *lot* of money on the table unspent. It is free to fail either on the upside with too many unspent assets, or on the downside, with spending ruin.

Perhaps this is well illustrated by the "one random run" spending line on Figure 9 for the 4% rule, representing the same series of returns as the one that was also used in Figures 5 and 6, earlier — so it is strictly comparable. Recall that in Figure 6 there was an early period of poor returns, followed by a flat period and then a recovery with better returns associate with more or less median spending. Under the 4% rule, the investor would run out of money on this particular run at about the 37th year, well short of our potential 45-year horizon. Why? Because in the early years, when the portfolio experienced a long period of disappointing returns, spending remained high — based on the 4% rule's built-in expectation that good and bad returns would offset each other over time. *That happens sometimes, but it doesn't always happen*. Here, spending is unable ever to recover because it was too high in the early poor return years, overly depleting the portfolio. A new spending level could be set after this period that would last the full remaining horizon, but it will have to be dramatically reduced below its earlier level, well below what it would have been had the spending been managed all along in an ARVA spending rule.

The upshot of the 4% rule, basically, is that one faces the opposing dangers of either spending too little while one could enjoy it, or spending too much and then either having to dramatically reduce spending in order to avoid running out of money — or actually running out of money!

Why does the 4% rule perform so poorly? The present value of the planned future 4% spending plan at no time bears any relation to the value of the portfolio, grossly violating the equality of the present value of future spending and the asset value — implying both a greater (unintended!) bequest plan than would have been desired if had been known, as well as a tolerance for a significant possibility of spending ruin. This spending rule certainly doesn't help to maximize the investor's consumption utility by minimizing spending risk for a given level of expected spending — if anything it maximizes the risk of spending too little or of spending too much, taking to little advantage of a growing portfolio and too much advantage of a declining portfolio. Maintaining that value identity allows rules like ARVA to maximize the sensible spend at every point in time, but it means that the spend itself will have volatility — volatility directly coming from the aggressiveness and volatility of the SAA investment policy.

Let's follow Bengen's more recent advice, and adjust the spending from time to time in our simulation, reflecting updated portfolio values as random returns evolve in our 12,000 runs of "history." His specific spending revision suggestions are a bit subjective and hard to model, and still not based on straightforward annuity math (unlike ARVA), so to explore the idea of adjusting spending on occasion, let's make up a mechancial rule designed to capture his intent of limiting over- and under-spending by making spending adjustments from time to time. We'll simply restart the spending process at the end of each 10 year period, spending 4% of the actual portfolio, increased by inflation, as the portfolio has evolved in each of the 12,000 runs of our simulation as of that date. We'll keep the inflation multipler consistent with the basic 4% rule: For example, after growing with inflation annually, as before, on the 10th year the new adjusted spend will be $4\% \times (1 + 2\%)^{10}$; we multiply the adjusted inflation times each of the 12,000 new portfolio values that period and that spending amount will remain constant other than that it will continue to increase with inflation until the next realignment with the actual portfolio values after another 10 years; then repeat. Figure 10 shows the resulting distribution of spending.

Figure 10: 4% Rule, Spending Over Time
With Spending Reset on Each 10th Year

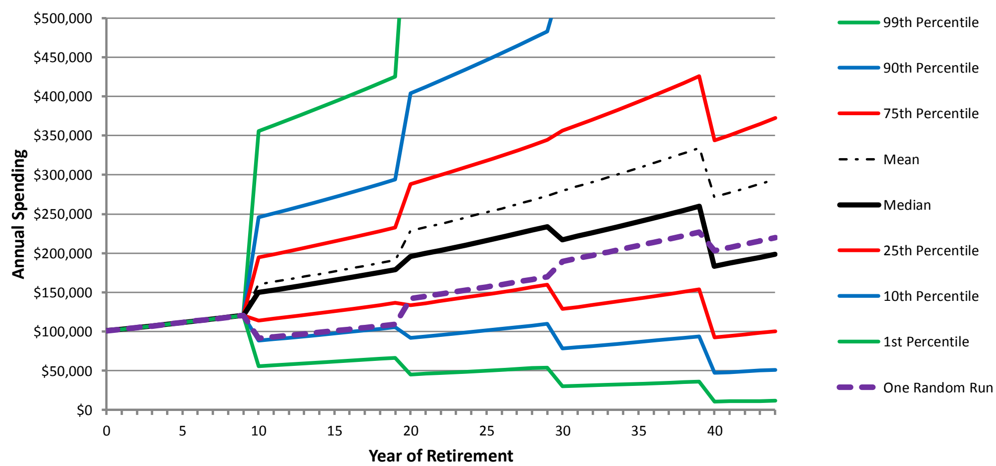

This procedure does ameliorate the problem quite a bit; spending tends a bit more to go up when the portfolio's value goes up, and down when it trends down. We now have a much broader distribution of spending levels, dependent on investment success, tending to look a bit more like Figure 6 than Figure 8, a good thing. The chance of complete ruin is much reduced, although the probability of unhappy levels of reduced spending is still quite large. Compare Figures 6 and 10 on the downside 1st and 10th percentiles; Figure 10's 1st percentile spending levels are much lower than Figure 6's, especially after about year 20. Smoothing, even with adjustments every 10 years, still provides a more risky consumption experience than ARVA.

Figure 11 shows the effect on the portfolio, demonstrating that as a result of the improvements in the spending distributions, the portfolio's value also has a somewhat improved distribution. There is less money left on the table unspent if times have been good; sadly, there is still a lot left on the table! And there is an improved, but still quite imperfect, tendency for the distribution to tighten and to reduce in value as the end of the horizon nears. A comparison of Figure 11 to Figure 5 still shows that the adjusted 4% rule still leaves too much money left in the portfolio unspent.

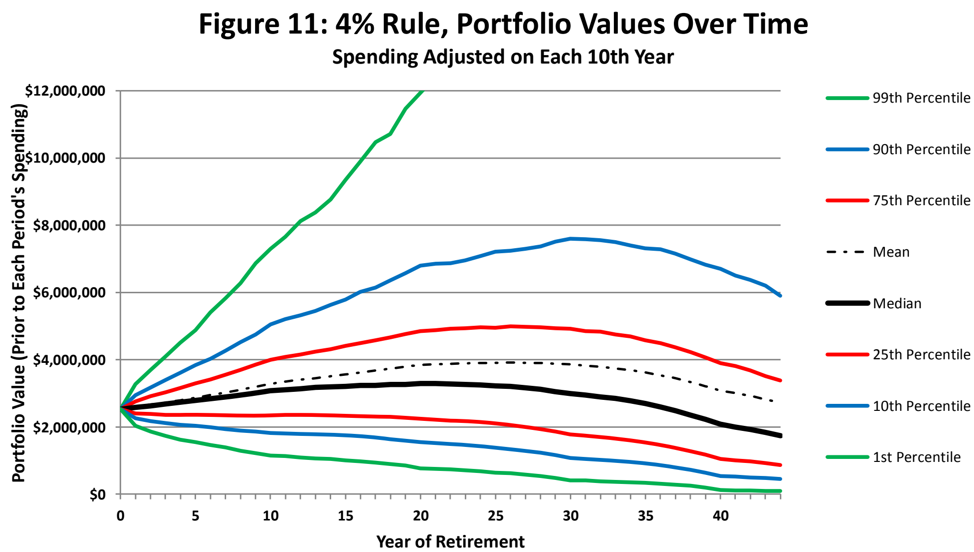

This is still not an efficient spending rule. The adjustments do improve things, but the real problem is that the identity of the present value of future spending to the portfolio value is not maintained. In effect, both the basic 4% spending rule and this adjusted version add risk to retirement spending — risk of either leaving too much money unspent, or worse, spending ruin. And worse yet, this is risk that could be eliminated without cost, simply by using an improved rule such as the ARVA rule.

For what it is worth, we aren't the first financial economists to criticize the 4% rule; Milevsky and Huang [2011] make their own criticism, and mention others prior to their own. But we believe that we are the first to demonstrate its failures graphically.

Thoughts on Smoothing in an Investment Context

Let's face facts: "Smoothing" is often sold as an appealing way to (a) hold an aggressive portfolio, but (b) not experience the risk from that portfolio. Unfortunately, this is a myth, perpetuated and popularized in investment contexts by the pension actuarial community, which continues to use it aggressively with defined benefit pensions to the detriment of pension benefit security. Its appeal lies in the apparent intellectual umbrella of mathematical assurances that ups and downs must average out over time; but the mathematical assurances are in error in the real world where market returns are not mean reverting. In general, it is good to remember that smoothing rules of any kind don't work well for de-risking processes that are independent and random from year to year, as we generally assume the markets to be — it is well known that risk to wealth increases over time under such conditions. The claimed benefits for smoothing are not true and can't be true.

Why? Here is a good example: When there are extended runs of poor performance, smoothed spending will overuse assets, necessarily causing either big reductions to future spending capability, or actual early ruin. Think about the period from 2000 through and including 2011, when the S&P 500 total return index grew only about a half percent per year on average for that 12 year period — spending a constant 4% during that period, when one is only making .5%, means 12 years of a portfolio shrinking at a compounded rate of 3.5%! Something will have to give. After just a few years, the portfolio would be so overdrawn as to force a very large and unpleasant spending reduction, far below expectation; smoothing is just not prudent.

If it isn't clear to every reader, *the original 4% rule, still widely used today, is simply a complete smoothing rule*, or polar case of such a rule, and is a very inefficient way to spend down one's portfolio. Adjusting it periodically reduces the inefficiency, but doesn't eliminate it. Smoothing does not remove risk; it *adds* risk, risk that is avoidable completely.

Our own ARVA spending rule, mentioned above, virtually or conceptually annuitizes the available funds each year to cover all remaining years of our planning horizon. It guarantees that the retiree will always continue to have some amount of income even in the presence of risky investments, and it leaves no money on the table if one lives to the end of the planning horizon. And it makes sense, because spending is always appropriate with respect to assets, time horizons, and interest rates, satisfying our value identity that the present value of future spending has to equal the value of the assets at all times. And, because it doesn't have to be biased to low spending in order to reduce the probability of ruin, it pays better overall whenever the strategic asset allocation choice is to take on any substantial amount of risk.

So we would seriously encourage advisors and retirees to consider annuity calculation-based spending rules such as ARVA, with frequent adjustments, in place of the 4% rule: one won't have unused assets if one lives through one's entire horizon, nor will one run out of money too early. Those risks can be eliminated, leaving only the pure investment risks to impact spending, as shown in Figure 6.

And if the regular adjustments are still too big to be comfortable under our suggested spending pattern, then perhaps one should adopt a more conservative SAA investment policy, shifting more from risky assets to hedging assets such as longer-term inflation-protected bonds! This, not smoothing, is the way to change one's risk exposure for real.

But regardless — what we've shown is what risk *is*, for a 4% rule smoothing investor.

Sensitivity of Results to the Simulation Model

As we said earlier, we are using a very simple and straightforward random walk model of returns as we simulate portfolio values and spending levels. This return model is in standard form, each period independent of the others, with the result being lognormal return distributions in each successive year.[11]

Nearly all analysts will accept this as a very good basic simulation model for demonstrating the points we want to make. It is not perfect, by any means, and there isn't any model that a consensus would agree to be "perfect." But it's perhaps the best and certainly the most widely-used general-purpose model. For those who would have preferred a model with fat tails, or a model with a bit of serial correlation, or a different shape to the spending rule, the pictures would be slightly different.

---

[11] See for one example of similar return simulation methods Chapter 12 of Hull [2009].

But the basic relationships between risk and return would be quite similar to what we are showing, given appropriately modest assumptions and assuming that any realistic model will exhibit increasing risk with increasing expected return; there would be substantial similarity in the lessons learned.

What we mean by this is that single-year asset-only return distributions will become increasingly wide as SAA risk increases; in particular, the distributions will become wider, that is, riskier, on the downside (and on the upside!).

And as a result, multi-period spending plan risk will always have a distribution of spending that gets wider and has both greater upside and greater downside, the greater the risk in the adopted strategic asset allocation policy.

Conclusion

It is easy for advisors and clients to think about returns. Risk is harder. And when it comes to evaluating the tradeoff between expected returns and risks, looking forward, most people — including most investment professionals — seem to get lost. Neither term, *expected return* nor *risk*, is used consistently and properly. The notion that an expected return is "the return that one should expect" has replaced on the street the correct reality that this is a statistician's use of the word "expected." A parent might say "I expect you to be home at 11 p.m." And the parent would expect that to happen with near certainty. But a statistician uses the term "expected" merely to indicate the average of a distribution, normally in investments a *wide* distribution, of expected outcomes or "realizations."

We've shown that in the widely used, asset-only single-period CAPM model, these realizations come from quite broad distributions, and that even if one expects, say, 6.9% per year in return for taking on some amount of risk, as we did above, the realized returns might be much lower or much higher than the expected 6.9%, where the likely deviation is in direct proportion to that risk. Recall how much taller we had to make Figure 2 than Figure 1, in order to show the effect of an 11% standard deviation on realizations! Risk is big, and should never just be skipped over in any SAA investment policy discussion!

And we've also shown that, in the more sophisticated multi-period CAPM, a world where assets are being used to fund a spending pattern, any risk taken on by the investor flows through to spending as the risk is realized — that is, it necessarily creates volatility in the spending pattern. Risk is taken in the hope that the spending pattern will be more favorable than what you would get from the riskless, hedged SAA investment policy, and it probably will be — but it might not. The risk is that the spending pattern will be less, maybe much less, than if the investor had simply hedged. The upside is that spending might be higher, and maybe much higher. But if the downside happens, will the investor be able to live with it?

And the risk doesn't go away with smoothing, as the famous and widely prevalent 4% rule attempts to do. Smoothing increases the risk of the extremes, of running out of money on the downside, as well as the risk of leaving substantial assets unused at the end of the planning period on the upside. And it tends to provide lower spending, on average, than ARVA would. The 4% rule adds risk, uncompensated risk, and shouldn't be used.

Here's the bottom line: If you take more strategic asset allocation policy risk, you might do much better either in single-period asset-only space, or in multi-period spending space. And on average, you can fairly expect to do better. But the thing is, that you might do a lot worse! You pays your money and you takes your chances. If you don't like the risk of doing worse, reduce the risk by adopting a more conservative strategic asset allocation policy.

We have shown you important pictures of what risk *is*, how it is shaped. The pictures should help you to better understand that more risk means a greater possibility of downside, not just a higher expected return. We'll repeat one last time: Don't like the risk? Change the asset allocation.

References

Bengen, William. 1994. "Determining Withdrawal Rates Using Historical Data." *Journal of Financial Planning* (March).

Bengen, William. 2012. "How Much is Enough?" *Financial Advisor* (May 1). Accessed on December 11, 2017 at https://www.fa-mag.com/news/how-much-is-enough-10496.html .

Charupat, Narat, Huaxiong Huang, and Moshe A. Milevsky. 2012. *Strategic Financial Planning Over the Lifecycle*. New York: Cambridge University Press.

Hull, John C. 2009. *Options, Futures, and other Derivatives.* 7th ed. Upper Saddle River NJ: Pearson/Prentice Hall.

Ibbotson, Roger G., and Rex A. Sinquefield. 1982. *Stocks, Bonds, Bills, and Inflation: The Past and The Future.* Charlottesville, VA: CFA Institute Research Foundation. Annually updated by Morningstar (to 2015) and Duff and Phelps (thereafter).

Ibbotson, Roger G., Moshe A. Milevsky, Peng Chen, and Kevin X. Zhu. 2007. *Lifetime Financial Advice: Human Capital, Asset Allocation, and Insurance.* Charlottesville, VA: CFA Institute Research Foundation, https://www.cfapubs.org/doi/pdf/10.2470/rf.v2007.n1.4580.

Lintner, John. 1965. "The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *The Review of Economics and Statistics*, Vol. 47, No. 1 (February), pp. 13-37.

Merton, Robert C. 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, Vol. 41, No. 5 (September), pp. 867-887.

Merton, Robert C. 1990. *Continuous-Time Finance*. Blackwell Publishing, Oxford, UK.

Merton, Robert C. 2014. "The Crisis in Retirement Planning." *Harvard Business Review* (July-August), https://hbr.org/2014/07/the-crisis-in-retirement-planning

Milevsky, Moshe A. and Huaxiong Huang. 2011. "Spending Retirement on Planet Vulcan: The Impact of Longevity Risk Aversion on Optimal Withdrawal Rates." *Financial Analysts Journal*, Vol. 67, No. 2 (March-April), pp. 45-58.

Mossin, Jan. 1966. "Equilibrium in a Capital Asset Market." *Econometrica*, Vol. 34, No. 4 (October), pp. 768-783.

Sexauer, Stephen C., and Laurence B. Siegel. 2013. "A Pension Promise to Oneself." *Financial Analysts Journal*, Vol. 69, No. 6 (November/December), pp. 13-32.

Sharpe, William F. 1964. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance*, Vol. 19, No. 3 (September), pp. 425-442.

Treynor, Jack L. 1962. "Toward a Theory of Market Value of Risky Assets." Originally unpublished. In *Asset Pricing and Portfolio Performance*, ed. Robert A. Korajczyk. London: Risk Books, 1999. This paper has been widely circulated since its original writing in 1962.

Waring, M. Barton and Duane Whitney. 2009. "An Asset–Liability Version of the Capital Asset Pricing Model with a Multi-Period Two-Fund Theorem." *Journal of Portfolio Management*, Vol. 35 No. 4 (Summer).

Waring, M. Barton and Laurence B. Siegel. 2015. "The Only Spending Rule Article You Will Ever Need." *Financial Analysts Journal*. Vol. 71 No. 1 (January-February).