

Retirement Planning Risk Tolerance

Posted in [Client Services](#)
Thursday, March 1, 2012

These are my reviews of two important articles on how to apply the risk tolerance concept to modeling client retirement capital needs and safe withdrawal rates.

"A Utility-Based Approach to Evaluating Investment Strategies"

by Joe Tomlinson

Journal of Financial Planning, February 2012

<http://www.fpanet.org/journal/AUtilityBasedApproach/>

Relevance: high

You already know how most retirement sufficiency analyses are conducted: you take a portfolio amount, assume a fixed time period for retirement, pick a distribution rate indexed to inflation, and then either look at actual stock/bond (or more complex portfolio mix) returns, or do a Monte Carlo analysis to evaluate the chances of failure.

What's missing? This very interesting article suggests that, to get a complete picture, we need to add a few more variables to the analysis. Start by assuming a variable lifespan. Instead of 20 or 30 years, have the model recognize that, just like stock returns, some clients will experience longer or shorter lives, and run the sufficiency analyses using actuarial tables and probability assumptions. Interestingly, the first thing that this one additional variable tells you is that all-bond portfolios are actually quite a bit riskier for people with variable and unpredictable lifespans (that is, those clients who happen to live in the real world), than for people who have an assumed fixed period of retirement.

That's the easy variable. A second variable is the AMOUNT of the shortfall or gain, on average, that the client will experience under each set of assumptions, which can be translated into the number of years the client will have to live without assets, or the amount of bequest that will be left to heirs. If a particular portfolio strategy results in many one- or two-year shortfalls, that may be qualitatively superior to a strategy that runs out of money less often, but whose shortfalls tend to span decades.

The hardest new variable, which is the main focus of this article, is a utility function: how does the client feel about these gains and shortfalls? If you look at all the outcomes of all the Monte Carlo runs for all the possible portfolios, you discover that many times the client leaves a large bequest to heirs at death, while other times the client is eating cat food and living under a cardboard box. Presumably, the client will be less concerned about the magnitude of the excess than about the possibility of a shortfall—but how much? How do you measure such a thing?

The article looks at the data from studies by Kahneman and Tversky, and conducts some independent research, which all tells us about what you would expect: 1) on average, people experience the pain of losses more keenly than the pleasure of gains; 2) this disparity is magnified when you talk about the pain of shortfall vs. the pleasure of excess in retirement; 3) individuals show a variation in this utility curve, so you cannot generalize about the client sitting in front of you; and 4) this stuff is more complicated even than that. As the excess goes up, the client will value it less; the client may be largely indifferent to the difference between a \$1 million and a \$2 million excess. Meanwhile, the difference between having just enough at death, on the one hand, and having to spend the final year of retirement eating cat food may be greater than, say, a similar one-year difference between living nine years vs. ten years on cat food. The utility functions will be curved and perhaps irregular, both above and below the graph, even if you can determine with some precision how much (or little) that person sitting in front of you experiences the pain of shortfall vs. the joy of leaving a bequest.

There are no conclusions here; the author is starting a discussion that will be picked up by others. He draws some utility curves, and shows you how to input them into an Excel spreadsheet, but suggests that further research can draw better ones. He offers some suggestions on how to run Monte Carlo models that incorporate variable life expectancy, magnitude of bequest or shortfall and utility function, and then invites others to try their own modeling techniques.

But you get a rough idea of how these additional variables can change your recommendations in the article's Table 2 and Table 3. The author starts with a base case of a portfolio 100% allocated to an inflation-indexed annuity whose distributions are assumed to be precisely sufficient for the retiree's future

lifestyle. Other portfolio combinations—various combinations of stocks and bonds, combinations of stocks and annuity—are compared favorably or unfavorably with the annuity option for persons with various aversions to shortfall.

For a person with a simple (and relatively modest) 2 to 1 experience of retirement shortfall pain over bequest pleasure, most traditional portfolios will work fine: for that person, the only portfolio that doesn't beat funding retirement with an inflation-indexed annuity is a 100% bond allocation. The same is roughly true for the person who has a 5 to 1 shortfall-to-bequest utility function.

But when clients move up to a 10 to 1 or 20 to 1—which the literature and the author's research suggest is more common—suddenly the optimal strategy is the 100% allocation to an annuity that provides retirement sufficiency certainty. The chances of shortfall—magnified by the client's greater aversion to them—are just too great for the client's comfort level.

Table 3 uses different return assumptions, and you quickly see that the optimal recommendation is highly-dependent on what you think the portfolio will return over the client's retirement period.

The goal here is not to move you to recommend more inflation-indexed annuities, but to start the dialogue about exploring client utility functions and incorporating two other important variables into the retirement analysis. Suddenly, reading this to the end, you realize that giving great advice about retirement distributions is way WAY more complicated than the current professional state of the art, or that the current research has yet explored. Expect to hear more about how to evaluate these utility functions, and maybe even a whole different kind of risk tolerance questionnaire, in the future. (p. 53)

"Spending Flexibility and Safe Withdrawal Rates"

by Michael Finke, Wade Pfau and Duncan Williams

Journal of Financial Planning, March 2012

<http://www.fpanet.org/journal/SpendingFlexibilityandSafeWithdrawalRates/>

Relevance: high

What's missing from the safe retirement withdrawal studies? Last month, in this same magazine, Joe Tomlinson had people smacking their foreheads when he introduced the idea of risk tolerance; that is, how willing are clients to risk running out of money in return for enjoying their retirement portfolio now? He used a utility function to define this willingness to accept the possibility of less in future years.

This article follows the same logic, giving retirees a risk aversion coefficient ranging from 1 to 10, and then looking at the implications for portfolio withdrawals and portfolio composition. It starts by defining the risks; a graph shows different stock/bond mixes ranging from 20% to 100% stocks, along with distribution rates ranging from 3% to 9%—and then shows the percentage chance, measured in number of years, when the clients would have to live on dog food. Based on this measure, the clients who are willing to risk a 9% distribution pattern (that is, 9% of the initial portfolio, taking out this dollar amount adjusted for inflation in subsequent years) have the best chance of success with a 100% stock portfolio. But the risk-averse client isn't even going to LOOK at that curve, which, at best, offers a 25% chance of impoverishment late in life. Further down, you can achieve what appears to be a 7% chance of running out of money with a 60% stock portfolio.

The goal, as Tomlinson pointed out, as these authors make clear, is to optimize the trade-off between spending more now and risking the chances of spending less later. If you plug in a risk aversion coefficient of 1, you are talking about a risk-taking client who is looking hard at that 9% distribution rate and ignoring the 3% option. On a graph, it takes all the curves on a different path, and you discover that the optimal tradeoff—the maximized certainty equivalence—can be found with a 70% stock allocation and a withdrawal rate of 7%. A close second is the 6% withdrawal rate with a 60% stock allocation, which is very nearly equal to a 100% stock allocation with an 8% withdrawal rate. This is living on the edge.

If the client sitting across the table has a risk aversion coefficient of 4, meaning she is slightly more timid but still above-average in the risk-taking world, you get a remarkably different picture: the optimal portfolio is 30% stocks and the optimal withdrawal rate is 4%. No other combination of withdrawal rates and allocations comes close; the 4% withdrawal rate is preferable for all portfolio mixes, and if the client decides to walk on the wild side and take a 5% distribution pattern, the optimal portfolio is between 40% and 50% stocks.

This, of course, is a remarkable difference between two clients who are on the upper end of the risk-taking scale, which tells you (something the authors don't point out) that this risk aversion coefficient tool is EXTREMELY sensitive, and if you get your client evaluation a teensy bit wrong, there is a high probability that the distribution path and portfolio will be more than a teensy bit off--a pattern advisors have learned to be wary of ever since the introduction of portfolio design optimizers back in the 1980s. The authors add in situations where there are fixed income sources that clients can draw on (Social Security? Defined Benefit Plan?), which of course raises their risk profile because at least some of the core living expenses are being met.

The authors state at the end that a client's willingness to take portfolio risk before retirement is equivalent to a willingness to accept shortfall risk after retirement, but I read through the article several times without finding any evidence that this is true. It's possible that the two are entirely different and not related at all, as FinaMetrica founder Geoff Davey has found, through thousands of questionnaires, that the willingness to risk life by skydiving or bungee jumping is generally unrelated to the willingness to assume downside risk in an investment portfolio. The skydiving adrenaline addict may be a wimp when it comes to a 10% stock market blip, and the person who can tolerate a 45% drop in portfolio value for the long-term higher returns may flinch at the thought of eating cat food after age 75.

That said, it only makes sense that some measure of risk tolerance would find its way into portfolio drawdown practices, just as the concept has long been a core part of pre-retirement portfolio planning. In two articles--Tomlinson's, and now this--we have come a very long way toward figuring out how to incorporate the risk tolerance concept into retirement planning. But we have also come hard against the reality that this is a very sensitive measure, and potentially very hard to determine with any precision. The next frontier in these studies is finding a tool that advisors can use, with some confidence, to ascertain a client's utility function or risk aversion coefficient. That, and perhaps we should eventually agree on one consistent thing to call it. (p. 44)

Trackbacks

No Trackbacks

Comments

No comments

The author does not allow comments to this entry