
The 4% Withdrawal Rule—Maybe Planners Have Been Wrong

By Joseph A. Tomlinson, CFP *Tue, Jul 6, 2010*

The complexity of the typical retiree's cash flows makes it nearly impossible for a planner to apply a simple rule of thumb like the 4% strategy, even if he or she wanted to.

Economists and financial planners often disagree, and one divisive issue involves the 4% rule for safe retirement withdrawals.

Every few months, an economist demonstrates that setting a retirement income strategy based on 4% inflation-adjusted withdrawals makes no sense. He or she usually argues that lifetime spending patterns should maximize utility, and that such patterns might look very different from an inflation-adjusted withdrawals strategy.

The financial planning community rarely comments on these arguments, perhaps because economists' papers are too technical for most planners. But, in this article, I'll attempt to do so. My goal is to explain the financial economists' arguments in a non-technical but in-depth way, and to respond to their criticisms from the perspective of a financial planner.

I've analyzed two financial economics papers on withdrawal strategies. The first paper, *The 4% Rule—At What Price?*, was published in three years ago by Jason Scott and Nobel laureate William Sharpe of Financial Engines, Inc., and John G. Watson of Stanford, but has recently received attention in the financial planning press.

The second paper, *Spending Retirement on Planet Vulcan: The Impact of Longevity Risk Aversion on Optimal Withdrawal Rates*, was published in March 2010 by academics Moshe Milevsky and Huaxiong Huang of the Individual Finance and Insurance Decisions Centre in Toronto. Both papers challenge the 4% rule, but from quite different perspectives. With apologies to the co-authors, I'll refer to these papers as Sharpe and Milevsky.

The 4% rule

First, let's define the 4% rule. Formulated by financial planner William P. Bengen in 1994, this guideline recommends a 30-year strategy of withdrawing 4% (pre-tax) of the value of the initial portfolio of investments each year, with annual increases by amounts based on actual inflation rates. Bengen assumed a portfolio of 50% to 75% stocks and back-tested the rule using overlapping 30-year time periods.

Since 1994, Bengen and other planners have produced variations on the 4% theme. They suggested that retirees could raise the initial withdrawal percentage if they agree not to increase withdrawals in years following future stock market declines. Such variations have been tested using both historic returns and Monte Carlo simulations.

The Sharpe paper

Financial economists often object that the 4% rule may produce consumption patterns that do not maximize lifetime utility. Sharpe uses this argument, and makes two other criticisms.

First, he calls the rule wasteful because it typically leaves excess funds at the end of life (he assumes a client with zero bequest motive). Using Monte Carlo simulations and assuming a 30-year retirement, his modeling predicts surpluses ranging from 10% to 20% of initial portfolio value.

The underlying math is straightforward. With a market-mimicking portfolio of stocks and bonds, and initial withdrawal rates low enough to hold the chance of failure to 0% to 5%, one would expect the average scenario to produce a surplus. To prevent a surplus would require spending more—and raising the probability of failure to a prohibitive 50%—or investing in low-return, risk-free assets. Sharpe's example produces a 4.46% initial withdrawal rate. He doesn't, however, advocate that planners recommend only risk-free assets. I'll have more to say about his overall conclusions later.

Sharpe also accuses the 4% rule of wasting 2% to 4% of initial assets that could be saved by using options strategies to eliminate sequence-of-returns risk. He makes the subtle point that financial markets don't compensate investors for sequence-of-returns risk, thus creating opportunities to use options as cost-free risk reduction. He uses a straight options strategy (buying and selling calls with different strike prices), as opposed to a life contingent strategy, which might include purchasing a variable annuity with a guaranteed lifetime withdrawal benefit or one of the newer standalone living benefit products that insure taxable separately managed accounts.

Using an efficient frontier graph with *Cost* on the horizontal axis and *Expected Utility* on the vertical axis, Sharpe demonstrates that the 4% rule produces an outcome falling well below the efficient frontier. He then shows how one might narrow the gap by spending the surplus and using options to eliminate unnecessary costs. These strategies do not move us all the way to the efficient frontier. To completely close the gap, he suggests shifting consumption between periods in ways that increase utility.

Sharpe also tests "glide-path" investment strategies, which systematically reduce portfolio volatility as a retiree ages. He generally concludes that such strategies are no less wasteful than constant mix strategies.

The question remains: What would Sharpe recommend to replace the 4% rule? Disappointingly, he doesn't provide details. Instead, he offers generalities such as, "There appears to be no doubt that a better approach can be found than that offered by combinations of desired constant real spending and risky investment," and "It is time to replace the 4% rule with approaches better grounded in fundamental economic analysis." We will need to wait for future papers for more specifics.

The Milevsky paper

Milevsky uses a different modeling approach in criticizing the 4% rule. His model uses stochastic mortality instead of a 30-year time horizon, and simplifies investments by assuming only risk-free assets. He provides calculations of utility-maximizing withdrawal rates using what economists refer to as “constant-relative-risk aversion” (CRRA) utility functions.

Milevsky’s modeling generates optimal withdrawal percentages, and his initial rates are in line with the 4% rule. The optimal withdrawals do not remain level, but instead decrease over the retiree’s lifetime. In one example, a 4.6% initial withdrawal rate at age 65 decreases slightly to 4.4% by age 75, but drops to 3.6% at age 90 and 2.2% at age 100. (Note, these percentages and all that follow refer to withdrawals as a percentage of the *initial* portfolio, and all withdrawals include adjustments for inflation.)

This pattern makes intuitive sense. Most 65-year-olds would think it reasonable to plan to spend roughly the same amount from ages 65 to 75, but they would be less likely to sacrifice income in the early years to ensure a similar annual income from ages 90 to 100—unless they are extremely averse to longevity risk.

That’s Milevsky’s point. He shows the initial withdrawal rates and the slope of planned withdrawals by age to be a function of the degree of longevity risk aversion. The risk-averse retiree is one who wants to make provision for a longer-than-expected life, so as risk aversion increases, initial withdrawal rates go down and the slope by age gets flatter. For example, the optimal initial withdrawal rate for a 65-year-old retiree might range from 4.1% to 6.3%, varying inversely with the degree of risk aversion.

Milevsky then shows the optimal withdrawal rates of a retiree who receives Social Security and/or other pension income. Other things being equal, more pension income will raise the optimal initial withdrawal percentage and increase the year-by-year tilt. For example, a 65-year-old retiree with \$1,000,000 in assets, a \$50,000 inflation-indexed pension, and low risk aversion might optimally consume 8% of the investment portfolio compared to 6.3% for the individual with no pension. For higher levels of risk aversion, similar proportional relationships apply, although the overall optimal withdrawal rates are lower and flatter by age.

Providing specific withdrawal percentages is not Milevsky’s goal, however. He’s out to prove that optimal planned withdrawal patterns decline by age, and may vary quite considerably as a function of risk aversion and pension income. In other words, 4% level withdrawals may be far off the mark in terms of utility maximization.

Though it’s not his central thrust, Milevsky also shows that investors can achieve higher withdrawal percentages by annuitizing some of their assets for life. A retiree with medium risk aversion and no annuitized assets, for example, would optimally plan to withdraw 4.6% at age 65, declining to 4.0% by age 80. At the opposite extreme, a client who annuitized 100% of assets could lock in a consumption rate of 6.3% for life.

Responding to the criticisms

Did Bill Bengen come up with a flawed idea in 1994, and have financial planners been doling out misinformation ever since? Should we tell clients to ignore the 4% rule and consult an economist? Or should we consider modifying the advice we dispense?

Before answering those questions, we should first ask, “How prevalent is the use of the 4% rule in actual financial planning practice?”

Stories in the *Journal of Financial Planning* or *Financial Planning Magazine* suggest that use of the 4% rule is pervasive. Hardly a month goes by without an article about the 4% rule or some variant of it. But I wonder how often the 4% rule is actually used in day-to-day financial planning practice.

Most clients have a myriad of cash flows with different timing. A client may have income streams that automatically adjust for inflation, like Social Security benefits, and others that do not, like corporate pensions. The client may have temporary expenses, such as a nearly paid mortgage, and one-time windfalls, perhaps from downsizing their housing. The complexity of the cash flows makes it nearly impossible for a planner to apply a simple rule of thumb like the 4% strategy, even if he or she wanted to.

This innate complexity forces most planners to customize their approaches to individual client cases. Ironically, the application of such customization may bring planning closer to the theoretical ideals of the financial economists mentioned here. For example, if a risk-averse client needs funding for basic recurring living expenses, the use of an immediate annuity would help overcome some of the objections to the 4% rule voiced by both Sharpe and Milevsky.

Different clients may have adequate income for basic living expenses, and prefer to allocate their savings for vacation spending. They may also want to spend their vacation dollars early in retirement rather than spread them over an uncertain future. This approach, which could be implemented without any knowledge of utility theory, would satisfy Milevsky’s utility-maximizing approach.

What can we learn from financial economics?

Indeed, planners can learn a lot from financial economists. For example, Sharpe illuminates the difficulties of trying to “*finance a constant, non-volatile spending plan using a volatile investment strategy.*” He also alerts us to the use of options strategies to mitigate sequence-of-returns risk. Although he doesn’t dwell on it in his paper, Sharpe also cautions against “free lunch” thinking. While it is feasible to reduce uncertainty on volatile investments, he says, it is not feasible to eliminate all risk and still earn a premium over the risk-free rate.

Milevsky establishes that the front-loading of spending may make sense for utility-maximizing spending plans that recognize longevity risk. His work also demonstrates how different degrees of risk aversion may affect optimal spending patterns and provides the numerical analysis to demonstrate the orders of magnitude of the differences. In addition, he shows how pension income may influence optimal spending and how the purchase of income annuities can increase funds available for retirement.

Unlike investment returns, utility cannot be measured precisely. Nonetheless, both of the papers discussed here argue convincingly that retirement planning should focus on the maximization of utility. To do that, we as planners will have to stretch our thinking beyond the application of mechanical rules.

Interestingly, both Sharpe and Milevsky assert that planners can use annuities to reduce their clients' investment risk and longevity risk. Our profession has focused on techniques aimed at making assets last a lifetime without adequately considering all the product alternatives. We need to listen to the criticisms of those outside our profession, and be more open to new ways of helping our clients plan for retirement.

©2010 RIJ Publishing LLC. All rights reserved.