THE 4% RULE—AT WHAT PRICE?

Jason S. Scott*, William F. Sharpe[a] *and John G. Watson*[b]

The 4% rule is the advice many retirees follow for managing spending and investing. We examine this rule's inefficiencies—the price paid for funding its unspent surpluses and the overpayments made to purchase its spending policy. We show that a typical rule allocates 10–20% of a retiree's initial wealth to surpluses and an additional 2–4% to overpayments. Further, we argue that even if retirees were to recoup these costs, the 4% rule's spending plan remains wasteful, since many retirees actually prefer a different, cheaper spending plan.

1 Introduction

Retirees must annually decide how to invest their retirement portfolios and how much to withdraw from them to cover their living expenses. Some retirees turn to financial planners for advice, whereas others consult brokers, investment publications, or web sites. Although these sources are quite different, their spending and investment advice is consistently the same—the 4% rule. This rule of thumb originated in the financial planning literature, and it was quickly adopted by many financial firms to advise their retail customers. Much of the financial press and many investor web sites now embrace the rule, and hence now it is the most endorsed, publicized, and parroted piece of advice that a retiree is likely to hear. Therefore, it behooves readers of this journal to be familiar with the rule's approach, features, and flaws.

A typical rule of thumb recommends that a retiree annually spend a fixed, real amount equal to 4% of his initial wealth, and rebalance the remainder of his money in a 60–40% mix of stocks and bonds throughout a 30-year retirement period. For example, a retiree with a $1MM portfolio should confidently spend a cost of living adjusted $40K a year for 30 years, independent of stock, bond, and inflation gyrations. Confidence in the plan is often expressed as the probability of its success, for example, in nine of ten scenarios, our retiree will sustain his spending. Modifications to this basic example include

*Corresponding author. Managing Director, Retiree Research Center at Financial Engines, Financial Engines, Inc., 1804 Embarcadero Road, Palo Alto, CA 94303, USA. Tel.: 650-565-4925; fax: 650-565-4905; e-mail: jscott@financialengines.com

[a]STANCO 25 Professor of Finance, Emeritus Graduate School of Business, Stanford University.

[b]Fellow, Retiree Research Center at Financial Engines, Financial Engines, Inc., 1804 Embarcadero Road, Palo Alto, CA 94303, USA.

changing the amount to withdraw, the length of the plan, the portfolio mix, the rebalancing frequency, or the confidence level. However, all these variations have a common theme—they attempt to finance a constant, non-volatile spending plan using a risky, volatile investment strategy. For simplicity, we refer to this entire class of retirement strategies as 4% rules, the sobriquet of its first and most popular example.

Supporting a constant spending plan using a volatile investment policy is fundamentally flawed. A retiree using a 4% rule faces spending shortfalls when risky investments underperform, may accumulate wasted surpluses when they outperform, and in any case, could likely purchase exactly the same spending distributions more cheaply. The objective of this study is to price these inefficiencies—we want to know how much money a retiree wastes by adopting a 4% rule. In the next sections, we review the 4% rule's history and examine its popularity. We then present a financial parable featuring two aging boomers and the single spin of a betting wheel. Our parable illustrates the flaws of the 4% rule, both qualitatively and quantitatively. Next, we use standard assumptions about capital markets and show that the approach of the 4% rule to spending and investing wastes a significant portion of a retiree's savings and is thus *prima facie* inefficient. Finally, we argue that a better solution can be obtained by formulating the retirement problem as one of maximizing the retiree's expected utility, an approach advocated by financial economists.

2   History

Not long ago, many financial planners estimated a retiree's annual spending budget using a mortgage calculator, an estimate of the average rate of return on the retiree's investments, and the retiree's horizon—the number of years that a retiree's investments had to support his spending.[1] Further, to include a cost of living increase, the planner would adjust the average nominal investment return downward by an estimate of the average inflation rate and compute the real spending. This method is only valid when all the future yearly investment returns and inflation rates are nearly equal to their estimated averages, and hence non-volatile.

First Larry Bierwirth (1994), and then William Bengen (1994) argued that since actual asset returns and inflation rates were historically quite volatile, retirement plans based on their averages were unrealistic. Bengen proposed an alternative strategy that retained the basic investment and spending strategies inherent in the mortgage calculation. In particular, he assumed that a retiree's assets were invested in a mix of stocks and bonds and annually rebalanced to fixed percentages. Further, he assumed that in terms of real dollars, a retiree's annual spending was constant and financed by a year-end, inflation-adjusted withdrawal from the portfolio. Hence, choosing a stock-bond mix and a withdrawal rate— the ratio of annual, real spending to initial wealth— specified a retirement plan. Now, for a given horizon, some of these plans would have historically performed better than all the other possibilities. Therefore, Bengen collected scenarios of past asset returns and inflation rates, simulated a number of plans under these scenarios, and identified the best performers.

Although a retiree wants the highest withdrawal rate possible, he also wants to sustain his spending throughout his retirement years. Bengen required that all his recommended plans be historically sustainable—the investments had to support all scheduled withdrawals for every historical scenario. Further, he sought the sustainable plans with the largest withdrawal rates. For example, using a portfolio mix consisting of between 50% and 75% stocks, Bengen (1994, 172) reported:

Assuming a minimum requirement of 30 years of portfolio longevity, a first-year withdrawal of 4 percent, followed by inflation-adjusted withdrawals in subsequent years, should be safe.

Larger (smaller) withdrawal rates were recommended for shorter (longer) horizons. However, the horizons of most interest were in the 30–40-year range and had withdrawal rates that were near 4%. For this reason, Bengen's approach is now commonly called the 4% rule.

In his original article, Bengen did not give a confidence level for his rule—he deemed it safe, since it had never historically failed. However, in a similar study, Cooley *et al.* (1998, 20, Table 3) reported a 95% historical success rate for a 30-year horizon, a 4% withdrawal rate, and 50–50% mix of stocks and bonds. This success rate increased to 98% when the percentage of stocks was increased to 75%. This study is often cited as the Trinity Study (all three authors are finance professors at Trinity University in San Antonio, Texas).

Bengen (1996, 1997, 2001, 2006a) extended his approach in a series of articles. For example, Bengen (1996) treats both tax-deferred and taxable accounts, and recommends that conservative clients reduce their risk with age by yearly decreasing their percentage of stocks by 1%. Bengen (1997) extends the investment choices to include Treasury Bills and small-cap stocks. For a summary of Bengen's work, we recommend his recent review (Bengen, 2006b). Cooley *et al.* (1999, 2001, 2003a) also continued to analyze withdrawal rates. Cooley *et al.* (1999) focuses on monthly versus annual withdrawals, and Cooley (2003a) concludes that investors "would benefit only modestly in the long run from international diversification."

When estimates of success rates are based on a small number of scenarios, they are prone to estimation error. This is particularly true for estimates that use overlapping historical scenarios. This problem led some investigators to develop market models—stochastic models of asset returns and inflation rate processes. A model's parameters are chosen so that the joint probability distribution of the processes reflects the average values, variances, and correlations commonly observed. An unlimited supply of scenarios can be numerically generated from these models, and then statistics, such as the success rate, can be computed using Monte Carlo methods. Further, in a few cases, estimates can be derived analytically.

George Pye (2000) simulated all-equity portfolios whose real returns were log-normally distributed with a mean return of 8% and a standard deviation of 18%. Pye concluded that his modified 4% rule would be safe for a 35-year horizon. Pye's strategy increases a portfolio's longevity by ratcheting down spending when markets perform poorly. We include his rule in the 4% class since its "focus is on sustaining the initial withdrawal" (Pye, 2000, 74) and invests in a volatile asset.

Most of the withdrawal rate research ignores a retiree's mortality and assumes a fixed planning horizon. An exception is the series of articles by Ho *et al.* (1994a, 1994b), Milevsky *et al.* (1997), and Milevsky and Robinson (2005). In particular, Milevsky and Robinson (2005) chose simple models for both the markets and mortality and developed estimates of success rates without using simulation. We recommend the article by Milevsky and Abaimova (2006) for a summary of this approach and its application to retirement planning.

Cooley *et al.* (2003b) compared the results obtained using historical data and market models, but their study "does not take sides on which methodology is better." Finally, in another approach, Spitzer *et al.* (2007) generate scenarios using a bootstrap algorithm to resample historical data with replacement.

3 Popularity

The authors conducted an informal survey of retiree guidance sources and their recommendations on spending and investment strategies. We were struck by the universal popularity of the 4% rule—retail

brokerage firms, mutual fund companies, retirement groups, investor groups, financial websites, and the popular financial press all recommend it. Sometimes the guidance explicitly references Bengen's work, the Trinity Study, or related research, but more often, it is presented as the perceived wisdom of unnamed experts. In this section, we report a sample of our findings.

Vice President of Financial Planning Rande Spiegelman of the Schwab Center for Financial Research wrote an article on retirement spending, subtitled "The 4% Solution", for the August 17, 2006 issue of *Schwab Investing Insights*®, a monthly publication for Schwab clients. In this article, he recommends the basic 4% rule that we used in the introduction to this study (30-year horizon, 4% withdrawal rate, 60–40% mix of stocks and bonds, and 90% confidence level).

T. Rowe Price's web site (2008a) suggests, "If you anticipate a retirement of approximately 30 years, consider withdrawing no more than 4% of your investment balance, pretax, in the first year of retirement. Each year thereafter, you'll want to increase that dollar amount 3% every year to maintain your purchasing power." A popular feature of this web site is the "Retirement Income Calculator" (T. Rowe Price, 2008b) that simulates 500 scenarios of returns for seven asset classes, where the monthly returns are assumed to be jointly normal. The calculator automatically accounts for minimum required distributions after age $70^1/_2$, attempts to decrease equity exposure every five years, and yearly inflates withdrawals by 3%. For a single retiree starting retirement at age 65, having a 30-year horizon, beginning with a $1MM portfolio, investing initially in a 60% equity mix (portfolio E), and withdrawing $3.3 K per month, the withdrawal rate is 3.96%, and the calculator predicts a 90% success rate.

The Vanguard Group (2008) advises "making withdrawals at rates no greater than 3% to 5% at the outset of your retirement..." They also provide a tool for retirees to determine how much they can annually withdraw in real dollars. The tool uses 81 historical scenarios, and "the monthly withdrawal amount shown by the tool is the highest level of spending in which 85% of these historical paths would have left you with a positive balance at the end of your chosen investment horizon." A retiree with a 30-year horizon is advised to withdraw at a rate of 3.75%, 4.75%, or 5.25%, if he is invested in a conservative (less than 35% equities), moderate (between 35% and 65% equities), or aggressive (higher than 65% equities) portfolio, respectively.

AARP (2008) indicates that "most experts" recommend 4% spending and that stocks should be added to the portfolio to "help your money last." Jane Bryant Quinn echoes Bengen's recommendations in the June 2006 *AARP Bulletin*. John Markese (2006), President of the American Association of Individual Investors (AAII), writes, "Most research and simulation studies conclude that portfolios with annual withdrawal rates of 4% or less and greater than 50% in stock are most likely to last throughout retirement."

Liz Weston (2008), writing for MSN Money Central, refers to Bengen and the T. Rowe Price Retirement Income Calculator (2008b) and recommends a 3–4% withdrawal rate for retirees with horizons of 30 or more years. Scott Burns (2004), business columnist for the *Dallas Morning News*, was an early champion of the Trinity Study. Also, Walter Updegrave (2007), *Money Magazine* senior editor, often fields questions on the 4% rule in his "Ask the Expert" forum. In his *CNN Money* article, "Retirement: The 4 Percent Solution", he recommends that a 65-year-old person withdraw money at the 4% rate, invest in a 50–50% mix of stocks and bonds, and shift gradually more to bonds with age. Further, he adds, "The reason many pros recommend this rate is because studies show that it

provides a high (but not absolute) level of assurance that your savings will last 30 years or more.

The explanation of the 4% rule's popularity is its simplicity—it is easy to understand and implement. Many of today's retirees use the rule and have benefited from its disciplined approach to spending and investment. However, as we shall see, that benefit comes with a price.

4 The parable of two boomers[2]

A simple example often yields insights that extend to more complicated venues—this is true in our case. In this section, we present a parable about two boomers, who share similar spending goals, but have radically different investment strategies. One of their strategies wastes money, and we will show that it can be replicated by a series of less-expensive strategies. In later sections, we use this same approach to show that the 4% rule wastes money.

Our two boomers, Eric and Mick, want to hear Bob Dylan in concert. Dylan's next concert has three-ticket tiers: floor pass ($125), general admission ($75), and pay-per-view ($50). Although both boomers would prefer a floor pass, unfortunately, each boomer has only $100 to spend. Fortunately, the concert promoters have set up a betting wheel and will allow Eric and Mick to each wager on a single spin. The betting wheel has three sectors of equal area that are equally likely to be selected. If a sector is selected, a chance on that sector will pay $1; otherwise the price of the chance is lost. Our boomers can purchase chances on any or all of the wheel's sectors. However, the promoters have contrived to make the costs of the three sectors differ[3]—they are 10 ¢, 30 ¢, and 40 ¢ for sectors one, two, and three, respectively, or in vector short hand (10 ¢, 30 ¢, and 40 ¢).

Eric decides to purchase 125 chances on each sector. He allocates his $100 wager as follows ($12.50, $37.50, and $50.00) to receive the payouts ($125,

$125, and $125). Eric's can't-lose, no-blues strategy is guaranteed to have a 25% return, and he will be enjoying the concert from the floor. However, Eric's friend Mick, who briefly attended the London School of Economics, likes high-return, high-risk gambles. Mick makes the purchases ($35.00, $15.00, and $50.00) and will receive the payouts ($350, $50, and $125). This strategy has a 75% expected return and a 127% standard deviation.

Mick has focused on his betting strategy and has lost sight of his quest, a concert ticket. In fact, if you take his goal into account, Mick is actually wasting money in three ways. First, Mick is paying for a surplus he is not planning to spend. If sector one occurs, Mick will get $350, buy a floor pass for $125, and be left with a $225 surplus. At the rate of a dime for a dollar, Mick is wasting $22.50 on the surplus, or alternatively, he has wagered only $77.50 of his $100 toward his ticket.

Second, Mick is also overpaying for his ticket chances—even if the surplus is eliminated. Mick's payouts without the surplus are ($125, $50, and $125). He has a two-third chance of a floor pass and a third chance of watching on pay-per-view. However, consider the alternative strategy that wagers ($12.50, $37.50, and $20.00) and costs just $70. It has the payouts ($125, $125, and $50) and would give Mick the same ticket chances as his current strategy and the same concert experience.[4] The alternative plan is less expensive because it pairs payouts and sectors so that the largest payouts correspond to the cheapest sectors.

Third, suppose that Mick ranks his ticket preferences using *Billboard* bullets: a floor-pass gets five bullets, a general admission ticket gets four bullets, and pay-per-view gets just one bullet. It is clear that Mick loves a live concert (he often performs himself). Further, he would prefer (as measured in expected bullets) a guaranteed general admission ticket (4 bullets) to a two-third chance at a floor pass and a third chance for pay-per-view ($3^2/_3$ bullets).

The betting strategy ($7.50, $22.50, and $30.00) would pay ($75, $75, and $75) and would guarantee a general admission ticket. Since Mick prefers the ticket chances of this $60 alternative strategy, he is wasting an additional $10 by following his current plan.

Our parable illustrates a number of key points. Eric wanted a floor pass to the concert, and he placed bets that paid $125 no matter what the wheel's outcome. Eric's fixed spending was paired with a risk-free strategy. On the other hand, Mick's wager has volatile payouts that overshoot the $125 goal when sector one is the outcome, but fall short for the remaining sectors. Generally, fixed spending is best supported by fixed payouts—volatile payouts lead to surpluses and deficits. In fact, Mick required $22.50 of his original stake to fund his surplus—wasted money. Further, Mick's particular strategy has wasted an additional $7.50 by overpaying for his spending distribution. Generally, when payouts are volatile, there may be alternative wagers that generate the same payout distribution, and some will be cheaper than others. Lastly, Mick arguably wasted another $10.00 by subscribing to an inferior ticket plan. More generally, when wagers and preferences are at odds, even more money can be wasted.

5   An experiment

Unfortunately, if a retiree adopts a 4% rule, he faces outcomes much like Mick's—he will waste money by purchasing surpluses, will overpay for his spending distribution, and may be saddled with an inferior spending plan. In this section, we describe a numerical experiment that demonstrates the first two problems. For simplicity, we limit our analysis to the specific, but representative, planning horizon of 30 years.[5]

Our experiment uses a simple market model—an economy with just two basic assets. The first is a one-year bond that pays a guaranteed annual 2% real return. A multiyear, risk-free bond can be replicated by purchasing a series of these one-year bonds. The second asset is a market portfolio consisting of all other bonds and all stocks, held in market proportions. The annual real returns on our market portfolio are independent and log-normally distributed with an expected value of 6% and a standard deviation of 12%. Because the returns are independent, they are serially uncorrelated. We note investors can create portfolios with any desired volatility by purchasing a mix of the two basic assets.

Our market model is similar to those used by investment consultants for asset allocation and asset liability studies. A 2% risk-free real rate is broadly consistent with the historical record for U.S. Treasury STRIPS and TIPS investment returns. In addition, our market portfolio assumptions imply a Sharpe ratio of a third, a fairly typical choice. While the actual market values of bonds and stocks vary over time, on average, bonds contribute about 40% of the value of the market portfolio and stocks 60%. Thus, a strategy that invests 100% in the market portfolio can be thought of as a 60% equity strategy.

An investor can guarantee a real dollar every year for 30 years by purchasing a series of zero-coupon, risk-free bonds. The cost of this investment is the sum of the discounted prices[6] $\$1/(1.02) + \$1/(1.02)^2 + \cdots + \$1/(1.02)^{30}$, which amounts to a little less than $22.40. Alternatively, if a retiree invests in a risk-free bond portfolio, he can safely withdraw at a yearly rate that is a bit more than $\$1.00/\$22.40 \approx 4.46\%$. This withdrawal rate—the guaranteed rate—is the maximum withdrawal rate that can be guaranteed to never fail. This risk-free strategy is analogous to Eric's strategy and is a special case of the 4% rule—the limit of zero investment volatility. This version of the 4% rule never has a surplus, never has a shortfall, and is the cheapest way to receive a constant, guaranteed payout every

year. If a cheaper investment were to exist, then there would be an arbitrage opportunity.

The risk-free strategy supports constant spending using the constant-return, risk-free bond—a perfect match. However, a typical 4% rule advises supporting constant spending using a volatile-return, stock, and bond mix—a mismatch. Our first set of experiments will analyze constant mix portfolios with six volatility levels, evenly spaced from 0% to 15%. These levels correspond to market exposures from 0% to 125%, or alternatively, equity exposures of from 0% to 75%.[7] We pair each of these portfolios with a range of withdrawal rates that brackets the guaranteed rate—4.0%, 4.25%, 4.46%, 4.75%, and 5.0%. Except for the pair that corresponds to Eric's strategy, our benchmark, these pairs will all exhibit at least one of Mick's problems.

We will report three metrics for each pair of investment and spending strategies: the failure rate, the cost of the surplus, and the overpayment for the spending distribution. A strategy's failure rate is the probability that the actual spending in the last year is less than the spending goal; in other words, the probability of a shortage. Failure rates (and success rates or confidence levels) have been the main focus of previous investigators, and we include them here for comparison. The final portfolio value, the surplus, is a function of random market returns. Hence, we can use the machinery of derivative pricing to find its fair price, the cost of funding the surplus. Similarly, we can compute the present value of the actual spending and the price of an alternative investment that delivers the same spending distribution, but at the cheapest price. The difference in these prices is the overpayment for the spending strategy. We refer the reader to Appendix A for more details.

For each pair of investment and spending strategies, we numerically simulated many equally probable, 30-year-long, scenarios. The portfolio value for each scenario was initially set to $100, without loss of generality, and the annual spending goal was set to the withdrawal rate times $100. We began each scenario by drawing a random real market return from its lognormal probability distribution, and recorded its value. Given the investment strategy and the returns on the risk-free bond and the market, we determined the real value of the portfolio at the end of the first year. Next, we computed the first year's actual real spending—the smaller of the portfolio's value and the spending goal—and deducted it from the portfolio. The actual spending amount and the portfolio's post-withdrawal value were recorded. This process was then repeated for the remaining 29 years. We then computed the metrics of our experiment using Monte Carlo estimates, for example, the failure rate was approximated by the percentage of scenarios for which there was a spending shortage. Again, we refer the reader to Appendix A for more details.

5.1   Failure rates

Table 1 summarizes the failure rates for our first set of experiments. The columns in Table 1 correspond to constant mix investment strategies (labeled with the strategy's volatility), and its rows to constant real spending strategies (labeled with the strategy's withdrawal rate). The investment portfolios are annually rebalanced to maintain a constant risk level. The numbers in the body of the table are the estimated failure rates, the percent of scenarios that fell short of the spending goal. For example, with 12% portfolio volatility and a 4.25% withdrawal rate, spending had a shortfall in 8.1% of cases. The standard errors of the estimates reported in this table are less than 0.01%. The failure rates reported in Table 1 are consistent with the previous studies.

Table 1 illustrates several key features of the 4% rule. First, the influence of portfolio risk depends critically on the spending level. For rates less than or equal to the guaranteed rate, adding risk necessarily introduces the possibility of failure. Moreover, the chance of failure increases with portfolio volatility.

Table 1 Failure rates for constant-mix portfolios.

Withdrawal rate	Constant-mix volatility					
	0%	3%	6%	9%	12%	15%
4.00%	0.0%	0.3%	1.9%	3.9%	5.7%	7.6%
4.25%	0.0%	1.9%	4.4%	6.3%	8.1%	9.9%
4.46%	0.0%	6.8%	7.9%	9.2%	10.6%	12.1%
4.75%	100.0%	22.5%	15.0%	14.0%	14.5%	15.4%
5.00%	100.0%	44.2%	23.4%	19.2%	18.4%	18.7%

Table 2 Cost of surpluses as a percentage of initial wealth for constant-mix portfolios.

Withdrawal rate	Constant-mix volatility					
	0%	3%	6%	9%	12%	15%
4.00%	10.4%	10.8%	13.0%	15.8%	18.8%	21.8%
4.25%	4.8%	6.3%	9.3%	12.5%	15.7%	19.0%
4.46%	0.0%	3.4%	6.8%	10.1%	13.5%	16.8%
4.75%	0.0%	1.2%	4.2%	7.5%	10.8%	14.2%
5.00%	0.0%	0.4%	2.7%	5.7%	8.9%	12.2%

However, for withdrawal rates above the guaranteed rate, investing in the risk-free asset is guaranteed to fail, so adding risk to the portfolio is the only possible hope for success. For example, by increasing the portfolio volatility from 0% to 3%, the 4.75% spending plan's failure rate drops from 100% to 22.5%. However, beyond a certain level, additional risk begins to increase the chances of failure. For example, with the 4.75% withdrawal rate, a portfolio volatility of 9% has the smallest failure rate, that is, 14.0%.[8] It is this pattern that led some researchers to devise strategies that minimize the failure rate and the risk of ruin.

5.2   Cost of surpluses

Retirees who follow a 4% rule will often generate portfolio surpluses that waste money. In Table 2, we report the results for our numeric experiment. The layout of Table 2 is the same as Table 1; however, the values in its body are estimated percentages of initial wealth that are spent on funding surpluses. The standard errors of all estimates are uniformly less than 0.05%. Clearly, withdrawing less (more) than the guaranteed rate and investing in the risk-free bond always (never) generates surpluses. However, for all rates, volatility adds significantly to the surplus and shifts money away from retirement spending. In particular, an investor withdrawing at the guaranteed rate (4.46%) and investing in the market portfolio (12% volatility) has allocated 13.5% of his portfolio to surpluses, and only 86.5% to actual spending. Strikingly, the surplus from this strategy could be used to increase *every* retirement payout by nearly 16% ($\approx$ 13.5%/86.5%).

5.3   Spending overpayments

In Table 3, we present our third metric—the spending overpayment. This table's structure is the same as its two predecessors, but the values in its body are

Table 3 Spending overpayments as a percentage of initial wealth for constant-mix portfolios.

Withdrawal rate	Constant-mix volatility					
	0%	3%	6%	9%	12%	15%
4.00%	0.0%	0.2%	1.1%	1.9%	2.5%	3.0%
4.25%	0.0%	0.7%	1.6%	2.4%	3.0%	3.5%
4.46%	0.0%	1.2%	2.1%	2.8%	3.4%	3.8%
4.75%	0.0%	1.7%	2.6%	3.3%	3.8%	4.2%
5.00%	0.0%	1.9%	2.9%	3.6%	4.1%	4.5%

the estimated percentages of initial wealth wasted by obtaining the spending distribution using the 4% rule versus a more cost-effective strategy. The standard errors of all estimates are uniformly less than 0.05%. The waste increases with increasing withdrawal rate, since relatively more money goes toward spending, and also increases with portfolio volatility. With a market portfolio investment (12% volatility), spending overpayments claim an additional 2–4% of a retiree's initial portfolio wealth. In the next section, we shall look at least cost strategies in greater detail and see how to eliminate overpayments.

6   Least cost spending strategies

We have identified and quantified the 4% rule's first two sources of waste—funding unspent surpluses and overpaying for spending. However, what investment strategy can a retiree use to recoup these losses? In this section, we answer that question and also demonstrate a useful diagnostic tool for discovering spending overpayments. For concreteness, we consider a retiree that has a 30-year horizon, starts with $100, withdraws at the 4.46% guaranteed rate, and invests in the market portfolio.

Figure 1 is a scatter plot of our retiree's last year of spending versus the cumulative market return over the 30-year period since his retirement began. For convenience, the cumulative market return has been annualized. Each point on the plot represents


Figure 1   Spending versus cumulative market return in year 30 for a 4.46% withdrawal rate and the market portfolio.

a scenario. As shown in Table 1, our retiree spent less than his $4.46 goal in 10.6% of the scenarios, and in fact, he spent nothing at all in 9.6% of them. The remaining 89.4% of the scenarios had sufficient funds to fully support the $4.46 spending goal. The Monte Carlo estimate of the cost of purchasing the last year's payout is 96¢. Note that if our retiree had invested in the risk-free bond, he would have been guaranteed a $4.46 payout in the last year. The cost of the bond investment would have been $4.46/$1.02^{30} \approx 2.46. Hence what happened to the $1.50 saved by allowing deficits? It was wasted on funding surpluses.

Our retiree spent 96¢ to fund the last year of his plan, but he could have obtained its spending distribution more cheaply.[9] To see why, consider points A and B in Figure 1. Point A is a scenario for which spending was $4.46 and the market averaged a mere 21 basis point increase over 30 years—times

Figure 2 Spending versus cumulative market return in year 30 for the least cost-spending strategy.

were tough. In contrast, point B is a scenario for which spending completely disappeared, and yet the market averaged a hefty 6.92% real, annualized increase. Since the cost of money is cheaper in good times than it is in bad times, our retiree could have saved money by finding an alternative investment that paid nothing in scenario A, $4.46 in scenario B, and kept the payouts in all the remaining scenarios unchanged. If our retiree exploited all these money-saving swaps, he would have a least cost strategy (Dybvig, 1988a, 1988b). Figure 2 is a scatter plot of the least cost-spending strategy versus the cumulative market return. The spending distributions of Figures 1 and 2 are the same—the spending amounts have the same values and frequency—however, the spending amounts have been reassigned to different market returns and have different underlying investment strategies. Our estimated price to implement the least cost strategy is 69 ¢, a savings of 27 ¢ over the 4% rule.[10]

Generally, scatter plots of spending versus the cumulative market return can identify overspending. A least cost strategy will plot as a smooth non-decreasing curve (see Figure 2), and spending will be higher when prices are cheaper. For example, the plot of constant, risk-free spending versus the market is a horizontal line, is non-decreasing, and is a least cost strategy. Inefficient strategies will either exhibit scatter (see Figure 1) or, if they are smooth, have a region where spending decreases with cumulative market return.

What investments should our retiree make to get the least cost? Because the shape of Figure 2 is approximately flat–ramp–flat, a financial engineer would suggest a strategy of buying and selling 30-year, European call options on the market portfolio. Our retiree would purchase calls with an exercise price equal to the rightmost value for which the payouts are zero. Simultaneously, he would sell an equal number of calls with an exercise price equal to the leftmost value for which the payouts equal $4.46. By adjusting the numbers of calls purchased and sold, a payoff diagram very similar to that shown in Figure 2 would be obtained. Specifically, suppose that N calls are bought at the strike price $K_b = F(P_0)$ and sold at $K_s = F(P_F)$, where P_0 and P_F (the failure rate) are the probabilities of a zero withdrawal and a shortage, respectively, and the function $F(\cdot)$ is the inverse cumulative probability distribution for the (log-normal) cumulative market return. The number of calls to buy and sell is $N = \delta \cdot W_0/(K_s - K_b)$, where δ is the withdrawal rate and W_0 is the initial portfolio value. The prices of the calls are easily determined from the Black–Scholes formula. For our example, $P_0 \approx 9.56\%$, $P_F \approx 10.58\%$, and $\delta \cdot W_0 = \$4.46$, and we get the values $K_b \approx \$2.12$ (2.53% annualized), $K_b \approx \$2.19$ (2.65% annualized), and $N = 57.97$. The net cost of the replicating call options is 68.5 ¢, which equals our Monte Carlo price to three digits.

The above analysis could be repeated to cover all 30 years of spending. Efficient strategies for each year would entail the purchase and sale of call options with different exercise prices. The total cost would equal the cost of all the options purchased minus the proceeds from selling all the options written. For this particular investment and spending combination, our retiree would save 3.4% of his initial portfolio by adopting the least cost strategy, where 0.27% of the savings can be attributed to the last year.

It is, of course, possible that no counterparty would offer the multiyear options necessary for a retiree

to replicate the 4% rule's spending distribution. If not, it is still possible to replicate the least cost strategy by following a dynamic strategy using the riskless bond and the market portfolio. Further, if enough retirees wanted the 4% rule's spending distribution and the least cost-investment strategy was sufficiently below other readily available investment options, then financial institutions would be motivated to offer new products and services that replicated the spending distribution at a cheaper price for the retiree, but at a profitable price for the institution.

The key to recouping surpluses is to just not fund them—just fund spending. The key to recouping overpayments is to spend relatively large amounts when payouts are cheap and to spend relatively small amounts when payouts are expensive. A retiree can get the 4% rule's spending distribution at a much cheaper cost by buying and selling call options or by implementing a dynamic strategy. However, does he really want this spending distribution or is there a cheaper distribution he would prefer?

7   Spending preferences and expected utility

Mick wagered $100 and had a two-third chance of winning a floor pass and a third chance of watching Dylan on pay-per-view ($3^2/_3$ bullets). Mick's gamble can be replicated for just $70 by eliminating a surplus and using a least cost strategy, but this is still wasteful, since Mick would have preferred a guaranteed general admission ticket (4 bullets), which he could have had by wagering just $60. In our parable, the money spent on surpluses and overpayments is only a lower bound on waste; choosing a strategy whose payouts are inconsistent with preferences wastes even more. In this section, we argue that this result is also true for a retiree who follows a 4% rule, but prefers a different spending plan.

Does a retiree really want the spending distribution of a 4% rule strategy? Consider a retiree who is planning on spending $40K a year for the next 30 years. Now, suppose that after the first year, his portfolio suffers a 20% loss. Without knowing much about our retiree, what would you expect him to do? Would he continue spending $40K, cut his spending by 20% to $32K, or do something in between? We might reasonably expect him to spread the pain of his loss across all his remaining years, and adopt one of the latter two options. However, if our retiree were a strict follower of the 4% rule, he will continue spending $40K a year, until his portfolio is exhausted.[11] It seems unlikely that a retiree would prefer fixed dependable spending in his early retirement years and the volatile spending of Figures 1 and 2 in his late retirement years. Our retiree could have addressed this issue much earlier, when he created his plan, either by investing in a riskless bond or by relaxing his requirement for fixed spending. However, without knowing more about his preferences, we do not know which option would have been more appropriate.

Financial economists advocate a very different approach toward retirement spending and investing. Building on the seminal work of Hakansson (1970), they suggest that a retiree's preferences concerning different amounts of income in different years be represented with a set of utility functions. Then they determine an integrated investment and spending policy that can be funded with the retiree's initial wealth and that will maximize the expected value of his overall utility.

Figure 3 places our earlier analysis in the context of retiree preferences and utility. Each point on the plot is a retirement strategy. The vertical axis plots the expected utility of a strategy, and the horizontal axis plots the cost of the strategy. The strategies that maximize the expected utility plot on the solid curve. The strategy labeled A represents the traditional 4% rule—an initial portfolio value of $100, withdrawal rate of 4.0%, and 60–40% stock-bond mix (i.e., the market portfolio). This strategy does not maximize expected utility—it has

Figure 3 Expected utility versus cost for various retirement strategies.

unspent surpluses and overpays for spending—and must plot below the curve. The least cost-spending strategy has the same spending distribution as the 4% rule, and therefore the same expected utility. However, this strategy recoups the $18.80 spent on surpluses (Table 2) and the $2.50 spent on overpayments (Table 3), and can be purchased for $78.70. The least cost-spending strategy (labeled B) plots to the left of the 4% rule strategy (labeled A) on Figure 3.

Generally, we do not expect strategy B to maximize a retiree's expected utility, and therefore, there may be many strategies that deliver the same benefit as A, but at a lower cost than B's. For example, suppose a retiree truly prefers constant spending—even a 1% chance of failure makes him extremely uncomfortable. He might be equally happy if he traded his current strategy that targets a $4 withdrawal, but has a chance of shortfalls, for a new strategy, which pays half as much, but is guaranteed. In Figure 3, we have labeled this case as C—it invests in the risk-free bond and costs $44.80. By knowing this retiree's preferences, we can make him just as happy as he was with the 4% rule, but at less than half the cost.

Of course, no advisor would recommend strategy B or C, since they both cost less than the available $100 and would leave money on the table. One strategy, labeled B′, fully funds the least cost-spending strategy, but uses the savings on surpluses and overpayments to increase all the payouts. This strategy costs $100, increases all payouts by more than 27%, and plots directly above the 4% rule in Figure 3. It is unlikely that this strategy would maximize the expected utility, unless our retiree enjoys the 4% rule's occasional spending shortages and his only mistake is failing to find the least cost way of achieving them. Hence, typically, there are many strategies with higher expected utility than B′ that cost the same. If C′ is one of these preferred strategies, it will plot above B′ in Figure 3. As an example, consider a retiree who has no tolerance for risk and prefers investing in the risk-free bond. He would use his full $100 to get a guaranteed $4.46 annual payout. This strategy's payout is at a minimum 10% higher than any payout of the 4% rule, and never fails to pay. It clearly dominates the 4% rule for all retirees, will not dominate the least cost strategy for a few retirees, and will be likely dominated itself for retirees who have a modest tolerance for risk.[12]

To actually reach the efficient curve on our figure, we need to know a retiree's utility function. However, much of the gap between the efficient curve and the 4% rule can be closed with a less formal description of a retiree's preferences. Basically, retirees need to understand that the cost of providing spending when markets perform poorly is often substantially greater than when they perform well.[13] If this point is well understood, then most retirees are likely to choose to spend less in bad scenarios in order to be able to spend more in good scenarios, versus spending the same amount in all scenarios. The willingness to make these spending tradeoffs and to assume the investment volatility needed to implement them will differ among investors, and therefore so should their retirement strategies.

8 Glide-path investment strategies

Our first set of experiments focused on constant-mix investment strategies—portfolios are annually rebalanced to maintain a constant volatility level. This was the original approach of Bengen (1994), and it is often recommended. Recently, however, *glide-path* investment strategies have experienced a surge in popularity. These strategies systematically reduce portfolio volatility as a retiree gets older, and several authors, including Bengen, have suggested that retirees adopt this alternative. Two years after his original article, Bengen (1996) concluded that, "All things considered, I recommend that you adopt a phase-down of one percent of your stock allocation each year . . ." Jennings and Reichenstein (2007) analyzed a number of lifecycle mutual funds and found that these funds, even when intended for post-retirement use, generally follow an investment glide path, averaging an equity percentage equal to approximately 120 less the target investor's age.

We have argued that the major flaw of the 4% rule is its attempt to support nonvolatile spending with volatile investing. Perhaps a glide path that systematically reduces risk could cause less inefficiency and reduce failure rates? Using our previous withdrawal rates and our previous volatility levels, we ran a numeric experiment to test this hypothesis. However, this time the volatility level was only used as the initial level for the first year's investments, and thereafter, the level was decreased each year so that in the 30th year the volatility was zero, and the portfolio was invested in just the risk-free bond.

The results of our glide-path experiments are reported in Tables 4–6, which correspond to Tables 1–3 of our constant-mix experiments. Comparing the failure rate tables, we notice that except possibly for the lowest volatilities, the glide-path rates are worse—even though the glide-path portfolios were on average less risky than the corresponding constant mix portfolios. The reason for this behavior is that glide-path strategies are just as likely as constant-mix strategies to have a series of poor early returns and run out of money. However, the glide-path strategy tends to lock-in these poor returns, decreasing the likelihood that future good returns allow the portfolio to recover and sustain spending.

A comparison of Tables 5 and 2, the cost of surpluses tables, shows that glide-path strategies do a bit better than constant-mix strategies in this category. Basically, a glide-path strategy generates a smaller surplus. For the market portfolio, the glide-path strategy spent approximately 2.5% less on surpluses than the corresponding constant-mix portfolio. However, a comparison of Tables 6 and 3, the overpayment tables, shows that the savings on surpluses are offset by higher overpayments. For the market portfolio, the glide-path strategy spent approximately 1–2% more on overpayments

Table 4 Failure rates for glide-path portfolios.

Withdrawal rate	Initial-mix volatility					
	0%	3%	6%	9%	12%	15%
4.00%	0.0%	0.2%	2.0%	4.1%	6.0%	7.7%
4.25%	0.0%	2.1%	5.3%	7.4%	9.1%	10.6%
4.46%	0.0%	9.8%	10.6%	11.5%	12.5%	13.5%
4.75%	100.0%	35.0%	21.6%	18.6%	17.9%	17.9%
5.00%	100.0%	63.9%	34.2%	26.2%	23.3%	22.3%

Table 5 Cost of surpluses as a percentage of initial wealth for glide-path portfolios.

Withdrawal rate	Initial-mix volatility					
	0%	3%	6%	9%	12%	15%
4.00%	10.4%	10.6%	11.9%	14.0%	16.3%	18.7%
4.25%	4.8%	5.7%	8.1%	10.6%	13.2%	15.8%
4.46%	0.0%	2.7%	5.4%	8.1%	10.8%	13.5%
4.75%	0.0%	0.7%	3.0%	5.6%	8.2%	10.9%
5.00%	0.0%	0.2%	1.6%	3.9%	6.4%	9.0%

Table 6 Spending overpayments as a percentage of initial wealth for glide-path portfolios.

Withdrawal rate	Initial-mix volatility					
	0%	3%	6%	9%	12%	15%
4.00%	0.0%	0.2%	1.5%	2.7%	3.7%	4.5%
4.25%	0.0%	0.9%	2.4%	3.5%	4.4%	5.1%
4.46%	0.0%	1.8%	3.1%	4.1%	5.0%	5.6%
4.75%	0.0%	2.5%	3.8%	4.8%	5.6%	6.2%
5.00%	0.0%	2.6%	4.2%	5.3%	6.0%	6.7%

than the corresponding constant mix strategy. In total, the surpluses and overpayments of glide-path strategies are comparable to constant mix strategies. Further, since we expect that the spending distributions generated by glide-path strategies are likely at odds with a retiree's true preferences, the costs of surpluses and overpayments are again only a lower bound on the total money wasted by adopting these strategies.

9   Conclusion

The 4% rule and its variants finance a constant, nonvolatile spending plan using a risky, volatile investment strategy. Two of the rule's inefficiencies—the price paid for funding its unspent surpluses and the overpayments for its spending distribution—apply to all retirees, independent of their preferences. For a typical rule, we used a market model to estimate that between 10% and 20% of a portfolio's initial wealth is being

allocated to surpluses, and an additional 2–4% is going toward overpayments. If the spending distribution of the 4% rule were inconsistent with a retiree's preferences, then the costs can be much higher. All in all, any retiree who adopts a 4% rule pays a high price.

Our approach can be easily extended to investigate other retirement rules of thumb and to use alternative market models. If a retirement plan generates unspent surpluses then our approach can price the surplus. A scatter plot of spending amount versus cumulative market return will quickly reveal whether a strategy is of least cost type. Strategies with overpayments will often generate a cloud of points (Figure 1), whereas least cost strategies will generate a non-decreasing curve (Figure 2).

Many practical issues remain to be addressed before advisors can hope to create individualized retirement financial plans that maximize expected utility

for investors with diverse circumstances, other sources of income, and preferences. While we still may be far away from such an ideal, 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. Despite its ubiquity, it is time to replace the 4% rule with approaches better grounded in fundamental economic analysis.

Appendix A. Pricing

Our simple two-asset economy is complete—any series of payouts, volatile or nonvolatile, can be dynamically replicated using just the risk-free bond and the market portfolio. Further, we can assign a fair, no-arbitrage price to the payouts using the economy's state price density function or pricing kernel (refer to Cochrane (2005) and Sharpe (2007) for a full discussion). This approach is often used to price derivative securities, and we use it in this study to price portfolio withdrawals for spending and surpluses. In this framework, the current price P of an asset that pays a random amount C_t in t years is equal to the expected value:

$$P = E[C_t \cdot M_t] \qquad (A.1)$$

where the random variable M_t is the pricing kernel. Generally, both C_t and M_t are functions of the random market returns $R_1, R_2, \ldots, R_t$, and the expected value $E[\cdot]$ in the above equation is with respect to the joint-density distribution of these returns.

Equation (A.1) can be used to price the surplus at the horizon $t = T$ and to price each annual spending withdrawal at years $t = 1, 2, \ldots, T$. The price of all withdrawals is equal to the sum of the prices for these annual withdrawals. Note that in our analysis, the total of all prices (withdrawals and surplus) will equal the initial portfolio value.

Usually, the pricing kernel M_t depends on the annual market returns via the cumulative market

return V_t. In this Appendix, we assume that all returns are gross returns (ratios of ending value to beginning value) so that in particular, the cumulative market return is just the product of consecutive annual market returns (i.e., $V_t = R_1 \cdot R_2 \cdots R_t$). Hence, whereas the payout often depends on the order of the returns, the pricing kernel's value is independent of order—a low return followed by a high return is equivalent to the high return followed by the low return. When the annual market returns have identical log-normal distributions and are independent, the pricing kernel is simply:

$$M_t = A^t / V_t^b \qquad (A.2)$$

The positive parameters A and b are given by the formulas:

$$A = (\sqrt{E_m \cdot R_f})^{b-1} \qquad (A.3a)$$

$$b = \frac{\ln(E_m / R_f)}{\ln(1 + S_m^2 / E_m^2)} \qquad (A.3b)$$

In the above equations, R_f is the total risk-free return, $E_m = E[R_t]$ is the yearly expected total market return, and $S_m = (\text{Var}[R_t])^{1/2}$ is the annual market volatility. For our experiments, we use the values $R_f = 1.02, E_m = 1.06$, and $S_m = 0.12$, and so we get the approximate values $A \approx 1.08$ and $b \approx 3.02$ for the two kernel parameters.

Because the exponent on the cumulative market return in Eq. (A.2) is positive, payouts are relatively cheap to secure when the market has high returns. This is a standard result of asset pricing theory: market prices must adjust until the total demand for claims paying off in states of scarcity will be less than that for claims paying off in times of plenty. To the extent that cumulative market returns serve as adequate proxies for overall consumption (scarcity or plenty), we expect the pricing kernel to be a monotonically decreasing function of cumulative market return.

Our Monte Carlo simulation generates a number of equally likely random paths, where each path is

a scenario of annual market returns of length T. For each path, we compute sample values for the payout and the pricing kernel. In particular, let $C_t^{(i)}$ and $M_t^{(i)}$ be the payout and kernel values for the ith path. Using N of these paths, we can estimate the expected value of Eq. (A.1) with a sample mean and get the following approximate price:

$$P \approx \frac{1}{N} \sum_{i=1}^{N} C_t^{(i)} \cdot M_t^{(i)} \qquad (A.4)$$

The error of this approximation can be made as small as we want by choosing the number of paths N to be sufficiently large.

For complete markets, Dybvig (1988a, 1988b) derived an elegant formula for computing the minimum price for receiving any payout distribution. Dybvig creates a dynamic investment strategy that replicates the payout distribution, but purchases payouts most efficiently—the highest payout is purchased at the cheapest price (when markets perform the best), the next highest is purchased at the next cheapest price, and so on. More formally, let $F_t(c) = P(\mathbf{C}_t \leq c)$ and $G_t(v) = P(\mathbf{V}_t \leq v)$ be the cumulative probability densities for the random payout $\mathbf{C}_t$ and market value $\mathbf{V}_t$, respectively. We can implicitly define the random payout $\mathbf{X}_t$ as a function of the cumulative market value $\mathbf{V}_t$ using the following equation:

$$F_t(\mathbf{X}_t) = G_t(\mathbf{V}_t) \qquad (A.5)$$

For any percentile, let us say the 75th, the 75th percentile of the market value will be paired with the 75th percentile of the payout distribution of $\mathbf{C}_t$, which becomes the 75th percentile of the payout $\mathbf{X}_t$. Hence, $\mathbf{X}_t$ increases with market value, decreases with pricing kernel, and has the same distribution as $\mathbf{C}_t$. Finally, if we replace $\mathbf{C}_t$ by $\mathbf{X}_t$ in Eq. (A.1), we obtain the formula for the minimum price.

The Monte Carlo simulation is easily modified to estimate the minimum price. After simulating N pairs of payouts and pricing kernels, each variable is separately sorted. The payouts are sorted in ascending order, the kernel values in descending order. The sorted values are re-paired and fed into Eq. (A.4). The result is an estimated minimum price. In practice, we generated a large number of paths ($N = 50,000$) for our estimates, and then further refined the results by averaging a large number (500) of independent estimates. For the results we present in this study, the standard deviation of the estimated prices was uniformly less than 5¢ for an initial wealth of $100.

Acknowledgments

Notes

[1] For brevity, we will refer to the typical retiree in the singular and use the male pronoun. However, all our arguments apply equally well to single females, a married couple, and partnerships.

[2] The author who is actually a member of the baby boom generation wrote this section. He assures his younger and older colleagues that the parable will have deeper meaning for readers of his cohort.

[3] The betting wheel's sectors have the same $1 payout, have the same chance of occurrence, but have different costs. Similarly, in an economy, a boom and a recession may have the same chance of occurring, but the cost of a $1 in the former state is typically much cheaper than in the latter state.

[4] We assume that Mick only cares about the distribution of tickets and not on the sector chosen by the wheel. More formally, Mick's utility is not state dependent.

[5] Modeling mortality is a very important aspect of retirement planning, and by assuming a fixed retirement period, we have made an extreme simplification—retirees drop dead at a specific, deterministic, future date. Further, we are implicitly assuming that any remaining portfolio value at that date is wasted. That said, any serious attempt to include mortality would introduce at a minimum annuities and bequest motives—two complexities that we defer

⁶ to another investigation. Annuities allow retirees to pool mortality risk, and bequests allow retirees to transfer surpluses to heirs.

[6] This is the formula for the present value of an ordinary simple annuity paying a dollar at the end of each year for 30 consecutive years.

[7] We included the leveraged market portfolio (125%) in the analysis to match the volatility of the oft-recommended 75% equity strategy. This implicitly assumes that investors can borrow at the risk-free rate. If borrowing costs are higher, the expected return for this portfolio can be adjusted downward.

[8] More simulations could be run to refine the estimate of the volatility that minimizes the failure rate for this, and any other, withdrawal rate.

[9] We assume that our retiree does not have state-dependent utility, cares only about the probability distribution of spending in each year, and does not care about his spending sequences.

[10] An alternative way to view the inefficiency with the payout plan shown in Figure 1 is to recognize that our retiree is exposed to two sources of risk. He does not know where on the horizontal axis the market's cumulative return will plot. This is market risk, for which an investor can expect to be rewarded in an efficient capital market. But even if the market's cumulative return were known, there may still be uncertainty concerning his payout, since there are multiple payouts for many of the vertical slices. Our retiree's payouts are not only dependent on the cumulative return on the market, but also on the order of the returns; that is, the market *path*. It is this path dependency that causes the observed multiplicity and creates additional risk. This additional risk, called *path risk* or *sequence risk*, is not rewarded in a capital market conforming to the standard assumptions of most asset pricing theories.

[11] Many devotees of the 4% rule recommend adjusting spending and investment after a dramatic market move, but some are strict adherents. For example, Guyton (2004) describes a retired couple, who after the bear market of 2000–2002, wonder if they "should reduce their withdrawals to keep their rate at roughly four percent" (of their depressed 2002 portfolio values). He advised them "that—on the contrary—it was fine if their current withdrawal rates approached six percent or even seven percent!"

[12] We can also plot Mick's gambles on Figure 3. The points A, B, and C, correspond to his original gamble ($100), the equivalent least cost gamble ($70), and the preferred guaranteed general admission ticket ($60), respectively. In Mick's case, the point B′ may correspond to a strategy

that has a third chance of a general admission ticket and a two-third chance of a floor pass versus his original gamble that has a third chance of pay-for-view and a two-third chance of a floor pass. For Mick, the point C′ would be Eric's gamble, the guaranteed floor pass, and since this gamble is optimal for Mick, C′ would plot on the curve.

[13] In Section 6, we estimated that the price of the least cost strategy for the 4% rule's spending distribution in year 30 was 69¢. This spending distribution pays $4.46 in over 89% of the top performing market scenarios. A risk-free bond would cover all market scenarios, pay $4.46, and cost $2.46. Hence, it costs another $1.97 to cover the bottom 11% of market scenarios; that is, almost three times as much to cover one-ninth of the scenarios.

References

AARP. (2008). "Managing Money in Retirement: Make Your Money Last." Availabe at: http://www.aarp.org/money/financial_planning/sessionseven/managing_money_in_retirement.html (accessed September 22, 2008).

Bengen, W. P. (1994). "Determining Withdrawal Rates Using Historical Data." *Journal of Financial Planning* **7**(4), 171–180.

Bengen, W. P. (1996). "Asset Allocation for a Lifetime." *Journal of Financial Planning* **9**(4), 58–67.

Bengen, W. P. (1997). "Conserving Client Portfolios during Retirement, Part III." *Journal of Financial Planning* **10**(6), 84–97.

Bengen, W. P. (2001). "Conserving Client Portfolios during Retirement, Part IV." *Journal of Financial Planning* **14**(5), 110–119.

Bengen, W. P. (2006a). "Baking a Withdrawal Plan 'Layer Cake' for Your Retirement Clients." *Journal of Financial Planning* **19**(8), 44–51.

Bengen, W. P. (2006b). "Sustainable Withdrawals." In: Evensky, H. and Katz, D. B. (eds.), *Retirement Income Redesigned: Master Plans for Distribution*. New York: Bloomberg Press.

Bierwirth, L. (1994). "Investing for Retirement: Using the Past to Model the Future." *Journal of Financial Planning* **7**(1), 14–24.

Burns, S. (2004). "Portfolio Survival: The Trinity Study." *Dallas Morning News*. December 7, 2004. Available at: http://www.dallasnews.com/sharedcontent/dws/bus/scottburns/readers/stories/SBportfoliosurvival.f5a90da.html (accessed September 22, 2008).

Cochrane, J. H. (2005). *Asset Pricing: Revised Edition*. Princeton, NJ: Princeton University Press.

Cooley, P. L., Hubbard, C. M. and Walz, D. T. (1998). "Retirement Savings: Choosing a Withdrawal Rate That Is Sustainable." *American Association of Individual Investors Journal* **20**(2), 16–21.

Cooley, P. L., Hubbard, C. M. and Walz, D. T. (1999). "Sustainable Withdrawal Rates from Your Retirement Portfolio." *Financial Counseling and Planning* **10**(1), 39–47.

Cooley, P. L., Hubbard, C. M. and Walz, D. T. (2001). "Withdrawing Money from Your Retirement Portfolio Without Going Broke." *CCH Journal of Retirement Planning* **4**(4), 35–41, 48.

Cooley, P. L., Hubbard, C. M. and Walz, D. T. (2003a). "Does International Diversification Increase the Sustainable Withdrawal Rates from Retirement Portfolios?" *Journal of Financial Planning* **16**(1), 74–80.

Cooley, P. L., Hubbard, C. M. and Walz, D. T. (2003b). "A Comparative Analysis of Retirement Portfolio Success Rates: Simulation Versus Overlapping Periods." *Financial Services Review* **12**(2), 115–129.

Dybvig, P. H. (1988a). "Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market." *Review of Financial Studies* **1**(1), 67–88.

Dybvig, P. H. (1988b). "Distributional Analysis of Portfolio Choice." *Journal of Business* **61**(3), 369–393.

Guyton, J. T. (2004). "Decision Rules and Portfolio Management for Retirees: Is the 'Safe' Initial Withdrawal Rate Too Safe?" *Journal of Financial Planning* **17**(10), 54–62.

Hakansson, N. H. (1970). "Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions." *Econometrica* **38**(5), 587–607.

Ho, K., Milevsky, M. A. and Robinson, C. (1994a). "How to Avoid Outliving Your Money." *Canadian Investment Review* **7**(3), 35–38.

Ho, K., Milevsky, M. A. and Robinson, C. (1994b). "Asset Allocation, Life Expectancy and Shortfall." *Financial Services Review* **3**(2), 109–126.

Jennings, W. W. and Reichenstein, W. (2007). "Choosing the Right Mix: Lessons From Life Cycle Funds." *American Association of Individual Investors Journal* **29**(1), 5–12.

Markese, J. (2006). "Developing a Withdrawal Strategy for Your Fund Portfolio." *American Association of Individual Investors Journal* **28**(11), 13–16.

Milevsky, M. A. and Abaimova, A. (2006). "Risk Management during Retirement. Retirement." In: Evensky, H. and Katz, D. B. (eds.), *Retirement Income Redesigned: Master Plans for Distribution*. New York: Bloomberg Press.

Milevsky, M. A., Ho, K. and Robinson, C. (1997). "Asset Allocation via the Conditional First Exit Time or How to Avoid Outliving Your Money." *Review of Quantitative Finance and Accounting* **9**(1), 53–70.

Milevsky, M. A. and Robinson, C. (2005). "A Sustainable Spending Rate without Simulation." *Financial Analyst Journal* **61**(6), 89–100.

Pye, G. B. (2000). "Sustainable Investment Withdrawals." *Journal of Portfolio Management* **26**(4), 73–83.

Quinn, J. B. (2006). "Get Smart with Your Money." *AARP Bulletin* **47**(6), (June). Reprinted and available at: http://bulletin.aarp.org/yourmoney/retirement/articles/simple_money_steps.html (accessed September 22, 2008).

Sharpe, W. F. (2007). *Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice*. Princeton, NJ: Princeton University Press.

Spiegelman, R. (2006). "Retirement Spending: The 4% Solution." *Schwab Investing Insights*®, August 17, 2006. Available at: http://www.schwab.com/public/schwab/research_strategies/market_insight/retirement_strategies/planning/retirement_spending_the_4_solution.html (accessed September 22, 2008).

Spitzer, J. J., Strieter, J. C. and Singh, S. (2007). "Guidelines for Withdrawal Rates and Portfolio Safety during Retirement." *Journal of Financial Planning* **20**(10), 52–59.

T. Rowe Price (2008a). "Investment Guidance and Tools: Retirement Planning: Living in Retirement: A Withdrawal Rule of Thumb." Available at: http://www.troweprice.com/common/index3/0,3011,lnp%253D10014%2526cg%253D12 210%2526pgid%253D7549,00.html (accessed September 22, 2008).

T. Rowe Price (2008b). "Investment Guidance and Tools: Retirement Planning: Retirement Income Calculator." Available at: http://www3.troweprice.com/ric/RIC/ (accessed September 22, 2008).

Updegrave, W. (2007). "Retirement: The 4 Percent Solution." *CNNMoney.com*. August 16, 2007. Available at: http://money.cnn.com/2007/08/13/pf/expert/expert.moneymag/index.htm (accessed September 22, 2008).

Vanguard Group. (2008). "Planning and Education: Retirement: Managing Your Retirement: Tap Your Assets: Determine How Much You Can Withdraw." Available at: https://personal.vanguard.com/us/planningeducation/retirement/PEdRetTapDetermineWDContent.jsp (accessed September 22, 2008).

Weston, L. P. (2008). "Make Your Money Last in Retirement; 5 Keys." Available at: http://articles.moneycentral.msn.com/RetirementandWills/RetireEarly/MakeYourMoneyLastInRetirement5keys.aspx (accessed September 22, 2008).