

Financial Planning Using Random Walks

John Norstad

j-norstad@northwestern.edu

<http://www.norstad.org>

February 3, 1999

Updated: November 3, 2011

Abstract

We develop an enhanced random walk model for retirement planning. Our model includes a new term to accommodate periodic additions to or withdrawals from a portfolio. We deal with the problems of modeling social security, salary growth, inflation, investment expenses, and asset allocation among cash, bonds, and stocks. We investigate a basic problem of post-retirement planning, the risk of outliving your money. We use a Monte Carlo simulation technique to implement the model in a computer program.

We assume that the reader is familiar with the lognormal random walk model as presented in reference [7].

Contents

1	Introduction	2
2	Developing the Model	3
2.1	Problem 1: How Much Money Do We Need?	3
2.2	Problem 2: How Do We Estimate the Parameters?	4
2.3	Problem 3: How Do We Enhance the Model?	7
2.4	Problem 4: How Do We Do the Computations?	8
3	Using the Model	10
3.1	Social Security	14
3.2	Rebalancing	14
3.3	Salary Growth	14
3.4	Asset Allocation and Risk/Return Tradeoffs	15
3.5	Investment Expenses	15
3.6	Assumptions and Possible Directions for Future Work	16
4	Post-Retirement Planning	18

List of Figures

1	15 Year Density and Cumulative Density Functions	11
2	Cumulative Density Function with Extra Monthly Savings	13
3	Post-Retirement Cumulative Density Function	19
4	30 Year Failure Rates	20

1 Introduction

In reference [7] we developed a lognormal random walk model for the S&P 500 stock market index. In this paper we build a similar model for a more difficult problem, planning for retirement in a realistic and complex situation similar to that faced by many people in the United States today.

To make it easier to illustrate how one develops a complex real-life random walk model, we'll take a concrete hypothetical example and develop all the details one step at a time.

Suppose Joe is 51 years old. He currently earns a gross salary of \$50,000 per year. He has a defined contribution retirement program at work where he contributes 5% of his gross salary via monthly payroll deduction and his employer adds a generous 2-for-1 matching contribution of 10% of his gross salary. Joe can contribute more than the standard 5% if he wishes (he doesn't currently), but he doesn't get any matching contribution from his employer for any such additional savings. Joe's retirement portfolio has been growing for many years and has a current market value of \$250,000. His portfolio has a fairly typical moderately aggressive asset allocation of 5% cash, 25% bonds, and 70% stocks.

Joe tells us that he would like to retire at age 66 and have enough money so that his standard of living in retirement will be the same as his standard of living in his last year of work prior to retirement. He wants his income during retirement to keep up with inflation. He expects to receive social security benefits, but other than that his only source of income will be his retirement portfolio savings from his job.

Joe wants to know what his chances are of achieving his goal. Does he need to save extra money? He has no clue how to answer these questions, so he has turned to us for help.

How might we go about trying to help Joe figure out his problem using a random walk model similar to the one we developed in [7]?

2 Developing the Model

We have four main problems to solve in developing a random walk model to help Joe plan for his retirement:

1. How much money does Joe need to accumulate in order to retire with his stated goals?
2. How do we estimate the parameters μ and σ for the random walk model?
3. How do we enhance the model to accommodate Joe's monthly savings and his employer's matching contribution?
4. How do we compute the density and cumulative density functions for the enhanced random walk model?

Let's solve these problems in order.

2.1 Problem 1: How Much Money Do We Need?

How much money does Joe need to accumulate in order to retire with his stated goals?

What does Joe mean by "the same standard of living?" Some expenses go up as we get older and some go down. Joe figures this will be a wash and decides he wants to have the same net dollar amount of income after retirement that he had before retirement.

We look at Joe's paycheck stubs and discover that he is paying the typical 7.5% of his gross salary in social security and Medicare taxes. He will not have to pay these taxes when he retires. In addition, Joe is currently putting away 5% of his gross salary as retirement savings, which is another expense he will not have when he retires. So Joe only needs to replace 87.5% of his last year's gross salary to maintain his net yearly income in retirement. Note that we do not subtract federal or state income taxes in this analysis, since Joe will still have to pay those taxes when he retires (we're assuming his portfolio is tax-deferred, as in a 401(k) or 403(b) plan.) For simplicity, we assume that Joe will be in the same tax bracket. This seems to be a reasonable assumption since his income will be about the same.

We next look at the most recent statement Joe has received from the Social Security Administration. This statement estimates that Joe's initial social security benefit will be 25% of his gross salary if he retires as planned at age 66, with a yearly cost of living adjustment (COLA) after he begins receiving benefits. Joe doesn't trust the politicians, and he believes that his actual benefits 15 years from now when he retires will likely be less than 25%. He asks us to use 20% as a conservative estimate to be on the safe side.

We have now determined that Joe's retirement portfolio must be large enough to replace 67.5% of his gross salary when he retires in order to maintain his standard of living.

Let $x = 67.5\%$ of Joe's gross salary in his last year of work prior to retirement. This is the amount Joe withdraws from his portfolio during his first year of retirement. In his second year he adjusts x for inflation and withdraws another x dollars. In his third year he adjusts x again, withdraws another x dollars, and so on throughout his retirement.

How big does Joe's retirement portfolio have to be to support these yearly inflation-adjusted withdrawals of x dollars? This is a complicated topic which we will discuss in detail in section 4. For now we make the assumption that 5% is a reasonable inflation-adjusted withdrawal rate.

Thus Joe needs to accumulate $1/5\% = 20$ times x dollars in his retirement portfolio. x is 67.5% of Joe's last year's gross salary. So Joe needs to accumulate 20 times 67.5% = 13.5 times his last year's gross salary.

This solves our first problem.

2.2 Problem 2: How Do We Estimate the Parameters?

How do we estimate the parameters μ and σ for the random walk model?

First we deal with salary raises and inflation. It may not be obvious what these issues have to do with parameter estimation, but that should become clear shortly.

Raises and inflation tend to be correlated. When inflation is high raises tend to be larger than when inflation is low. Thus it makes sense to deal with the two issues at the same time. We ask Joe to supply a history of his raises over the last 10 years. We compare his raises to the time series data for the Consumer Price Index (CPI) over the last 10 years. We do some quick arithmetic to determine that Joe's average raise in excess of inflation over the last 10 years has been 1.5%. For our model, we'll assume that this pattern continues over the next 15 years.

The easiest way to deal with inflation and raises is to change our unit of measurement for Joe's portfolio value from dollars to multiple of Joe's current salary. For example, today Joe's salary is \$50,000 and his portfolio value is \$250,000. In our new units of measurement Joe's salary is 1.0 and his portfolio value is 5.0. Let's suppose that over the next year the CPI increases by 2.5%, so Joe's salary increases by $2.5\% + 1.5\% = 4.0\%$. Let's also suppose that Joe's portfolio value increases by 20% due to the combined effects of his monthly savings, his employer's matching contribution, and a good year for the markets.

After one year, Joe's salary grows by 4% to \$52,000 and his portfolio grows by 20% to \$300,000. In our new units of measurement, Joe's salary is still 1.0 (by definition), and his portfolio value is $\$300,000/\$52,000 = 5.7692$. In our old units of measurement Joe's portfolio grew by 20%. In our new units of measurement, however, Joe's portfolio has grown by only $(5.7692 - 5.0)/5.0 = 15.38\%$.

In this example, the 20% return is called a "nominal return" and the 15.38% return is called a "real return relative to Joe's salary growth." The formula for computing the real return as a function of nominal return and salary growth rate is:

$$\text{real return} = \frac{\text{nominal return} - \text{salary growth rate}}{1 + \text{salary growth rate}}$$

You can confirm this by plugging in the numbers from our example:

$$\text{real return} = \frac{20\% - 4\%}{1 + 4\%} = \frac{.16}{1.04} = 15.38\%$$

This analysis suggests that a good way to model inflation and salary growth is to use real rates of return relative to Joe's salary growth rather than nominal returns. So when we use our time series data to estimate the model parameters we need to apply the conversion formula given above.

The next issue that affects parameter estimation is investment expenses.

In our S&P 500 model we made the assumption that we could invest money in the model without having to pay any kind of fee. This assumption is unrealistic with real-life investing. Brokers, mutual fund investment companies, insurance companies, and retirement plan management companies all charge fees. These fees are usually set as a yearly percentage of the portfolio value and are measured in "basis points", where one basis point is 0.01%. We do a bit of research and discover that Joe's investment expenses in his retirement plan are 35 basis points (0.35%).

To properly model Joe's portfolio, when we do our parameter estimation, before converting from nominal to real returns, we must subtract 0.35% for Joe's investment expenses.

To summarize, Joe's portfolio is 5% cash, 25% bonds, and 70% stocks. His yearly salary increase is 1.5% in excess of inflation. His investment expenses are 35 basis points.

We're now ready to do the parameter estimation for our random walk model. We have the following four sets of historical market time series data from 1926 through 1994:¹

C = Cash = 90 day US Treasury bills yearly returns
 B = Bonds = 20 year US Treasury bonds yearly returns
 S = Stocks = S&P 500 index yearly returns
 I = Inflation = Consumer Price Index percent change

We construct a new time series N for the nominal return of Joe's portfolio after expenses as follows:

$$N = 0.05 \times C + 0.25 \times B + 0.70 \times S - 0.0035$$

We then convert from nominal returns to real returns relative to Joe's salary growth:

$$R = \frac{N - (I + 0.015)}{1 + (I + 0.015)}$$

We finally estimate the parameters μ and σ for our model in the same way we did for the S&P 500 model, by converting to continuous compounding and finding the mean and standard deviation:

$$\begin{aligned} \mu &= E(\log(1 + R)) = 0.037405 \\ \sigma &= \sqrt{\text{Var}(\log(1 + R))} = 0.146836 \end{aligned}$$

Note that the expected continuously compounded real rate of return after expenses for Joe's portfolio relative to his salary growth is only 3.74%, with a rather large standard deviation of 14.68%.

¹The data is from Table 2-4 in reference [2].

2.3 Problem 3: How Do We Enhance the Model?

The random walk model we developed in [7] is:

$$\frac{ds}{s} = e^{\mu dt + \sigma dX} - 1 \quad \text{where } dX \text{ is } N[0, dt]$$

In this stochastic differential equation, dt is an infinitesimal time interval, ds is the change in the portfolio value over the time interval, and dX is a random variable with variance dt .

This equation has a useful property for computation. dt does not have to be an infinitesimal time interval. It can be any time interval! For example, the equation works fine with $dt =$ one day, $dt =$ one month, or $dt =$ one year. We change notation to recognize this:

$$\frac{\Delta s}{s} = e^{\mu \Delta t + \sigma \Delta X} - 1 \quad \text{where } \Delta X \text{ is } N[0, \Delta t]$$

Unfortunately, this model does not deal with periodic savings or withdrawals from the portfolio. We have to add an extra term to accommodate this:

$$\frac{\Delta s}{s} = e^{\mu \Delta t + \sigma \Delta X} - 1 + \frac{k \Delta t}{s} \quad \text{where } \Delta X \text{ is } N[0, \Delta t] \quad (1)$$

$$s(t + \Delta t) = s(t)e^{\mu \Delta t + \sigma \Delta X} + k \Delta t \quad (2)$$

where k is the constant yearly extra amount added to ($k > 0$) or withdrawn from ($k < 0$) the portfolio in even installments at the end of each time period Δt .

Equations (1) and (2) are our enhanced random walk model for retirement planning.

In Joe's case he contributes 5% of his gross salary via monthly payroll deduction and his employer contributes a 10% matching contribution. Thus for Joe's model we use $k = 0.15$ and $\Delta t = 1/12 = 0.08333 \dots$

We now have a complete random walk model for Joe's retirement plan.

To summarize everything we have done so far, the model is given by equations (1) and (2) above with the following parameters:

$$\begin{aligned} \mu &= 0.037405 \\ \sigma &= 0.146836 \\ s_0 &= 5.0 \\ k &= 0.15 \\ \Delta t &= 1/12 = 0.8333 \dots \\ t &= 15 \end{aligned}$$

Joe's goal is to accumulate a total of 13.5 times his salary, $s(t) = 13.5$.

Joe would like to see the cumulative density function for $s(t)$. This would give him an answer to his primary question “what are my chances of reaching my goal?”

This leads us to our last problem number 4.

2.4 Problem 4: How Do We Do the Computations?

How do we compute the density and cumulative density functions for the enhanced random walk model?

In the S&P 500 model we developed in [7] this was an easy problem because the ending value $s(t)$ was a simple lognormally distributed random variable, so we were able to use standard built-in functions and readily available computer code to compute the functions (see section 6.5 in [7]). We’ll see that this is unfortunately no longer the case with our enhanced model, so we’re going to need to do some more work.

Let $n = t/\Delta t$ and $\Delta X_i =$ independent $N[0, \Delta t]$ random variables for $i = 1 \dots n$. It is a bit laborious but quite easy to show by induction that:

$$s(t) = s_0 \prod_{i=1}^n e^{\mu\Delta t + \sigma\Delta X_i} + \sum_{i=1}^{n-1} k\Delta t \prod_{j=i+1}^n e^{\mu\Delta t + \sigma\Delta X_j} + k\Delta t \quad (3)$$

The first term is the growth of our initial portfolio value over the full t year period.

The second term is the sum of the growth of each monthly savings amount $k\Delta t$ over the months from the time the savings are added to the portfolio through the end of the full period.

The last term is the last month’s savings, made at the end of the month.

We can rewrite this equation as:

$$s(t) = s_0 e^{\mu t + \sigma \left(\sum_{i=1}^n \Delta X_i \right)} + k\Delta t \sum_{i=1}^{n-1} e^{\mu(n-i)\Delta t + \sigma \left(\sum_{j=i+1}^n \Delta X_j \right)} + k\Delta t$$

Can we simplify this equation further? Note that $s(t)$ is a sum of n lognormally distributed random variables plus a constant, but the random variables are not independent, and even if they were, their sum would not be lognormally distributed. (The product of two independent lognormally distributed random variables is lognormally distributed, but their sum is not.)

We seem to be at an impasse. We do not have anything corresponding to the simple equations we presented in section 6.5 of [7] for our enhanced model. $s(t)$ no longer has a simple distribution which we can easily compute, so our retirement savings graphs are going to be harder to calculate than were our graphs for the S&P 500 model.

How might we go about computing the density and cumulative density functions for $s(t)$?

We might first try to derive the equation for the cumulative density function from the equation (3) given for $s(t)$ above. This is certainly possible. If we do it, we end up with an n -dimensional definite integral which must be evaluated over a complex region in n -space. We could try to solve this integral using calculus, or we could try to use numerical integration techniques to evaluate it. While perhaps feasible, these approaches are certainly complicated and far from obvious. At least, they are far from obvious to your ignorant author, who is much too lazy to follow this line of thought further if there is any reasonable alternative.

Fortunately, there is indeed a reasonable alternative. We can fall back on the programmer's motto: "When all else fails, use brute force."

Our goal is to graph the density function and the cumulative density function for $s(t)$. Suppose we simulate some large number of random walks in our enhanced model, say 100,000 of them, and record all the ending values. We end up with a long list of 100,000 ending values.

To graph the density function, we partition the range of ending values into a number of evenly spaced "buckets" and count up the number of ending values in each bucket. We draw a histogram of the result. If our buckets are sufficiently narrow, say only one screen pixel wide in our graph on the screen, and if our number of samples is sufficiently large, the result is a close approximation to the density function graph. The graph is a bit "jagged" due to the approximation, but it's more than good enough to see its shape and essential characteristics.

The cumulative density function is also easy to compute. To estimate $\text{Prob}(s(t) < k)$ for any k , simply walk through the list and count up the number of elements which are less than k . The function value is the count divided by 100,000. Graphing the cumulative density function is also easy, and it isn't even all that "jagged!"

This brute force technique is called a "Monte Carlo simulation." It is the technique we have implemented for our enhanced model. Fortunately, modern personal computers are so fast that this technique is quite feasible and works well in practice.

We're making enormous progress on Joe's retirement planning problem. We can now graph the functions and answer his question "what are my chances?" We show the results in the next section.

3 Using the Model

Figure 1 shows the density and cumulative density functions for $s(15)$ for Joe's current retirement savings.²

From the cumulative density function, we see that the probability that Joe will meet his goal of $s(15) = 13.5$ is only about 40%. The median value (50/50 chance) of $s(15)$ is a bit less than 12, well short of Joe's goal by more than 1.5 times his last year's gross salary.

Joe decides that he needs to save more money to improve his chances of meeting his goal. He'd like to at least improve his odds to 50/50. How much extra does he need to save?

Recall that in the beginning we assumed that Joe was saving 5% of his gross salary. Together with information about Joe's withholding for social security and Medicare taxes (7.5%) and Joe's estimate of his future social security benefits (20%), we computed that Joe's retirement portfolio needed to be large enough to replace $100\% - 7.5\% - 20\% - 5\% = 67.5\%$ of his last year's salary. We assumed a 5% inflation-adjusted post-retirement withdrawal rate, so we multiplied this number by 20 to get our target number of 13.5 times Joe's salary.

Now we are going to change the amount Joe saves, so we need to parameterize this calculation. Let

$x =$ Joe's total savings as a percentage of his gross salary, not counting the fixed 10% matching contribution from his employer.

Note that when we add back in the matching contribution we get:

$$k = x + 0.1$$

Our new accumulation target for when Joe retires at age 66 is:

$$\begin{aligned} \text{Target } s(t) &= 20(100\% - 7.5\% - 20\% - x) \\ &= 20(0.725 - x) \\ &= 14.5 - 20x \end{aligned}$$

We want to find the value of x which makes the median value of $s(t)$ equal to this target.

²All the graphs in this paper were drawn using the "Random Walker" program [6], available at the author's web site.




Figure 1: 15 Year Density and Cumulative Density Functions

There's no easy way to compute the exact median value of $s(t)$, but we can get a value that's close enough for our purposes by setting $\sigma = 0$ in equation (3) we derived for $s(t)$ on page 8 above:

$$\begin{aligned}
 \text{Median } s(t) &= s_0 e^{\mu t} + \sum_{i=1}^{n-1} k \Delta t e^{\mu(n-i)\Delta t} + k \Delta t \\
 &= s_0 e^{\mu t} + k \Delta t \left(\sum_{i=1}^{n-1} (e^{\mu \Delta t})^{n-i} + 1 \right) \\
 &= s_0 e^{\mu t} + k \Delta t \sum_{j=0}^{n-1} (e^{\mu \Delta t})^j \\
 &= s_0 e^{\mu t} + k \Delta t \frac{(e^{\mu \Delta t})^n - 1}{e^{\mu \Delta t} - 1} \\
 &= s_0 e^{\mu t} + k \Delta t \frac{e^{\mu t} - 1}{e^{\mu \Delta t} - 1}
 \end{aligned}$$

Set Target $s(t) = \text{Median } s(t)$, substitute $k = x + 0.1$, and solve for x :

$$x = \frac{(14.5 - s_0 e^{\mu t})(e^{\mu \Delta t} - 1) - 0.1 \Delta t (e^{\mu t} - 1)}{20(e^{\mu \Delta t} - 1) - \Delta t (e^{\mu t} - 1)}$$

Use a calculator, spreadsheet program, or some other tool to evaluate this expression for the values of our parameters. The result is 9.30%.

Thus in order to improve his chances of meeting his goal to 50/50, Joe needs to save a total of 9.30% of his gross salary. He's already saving 5%, so he needs to save an extra 4.3%. His current gross salary is \$50,000, so this is extra savings of \$2,150 per year or \$179.17 per month.

We now reset the parameter k in our model to Joe's new 9.30% savings plus his employer's 10% matching contribution = 19.30% = 0.193. Joe's new savings target is $14.5 - 20 * 9.3\% = 12.64$. We use our program to draw a new cumulative density graph. Figure 2 shows the graph.

As expected, the median ending value is now about 12.64, Joe's new target.

Joe now has a 50/50 chance of meeting his goal. But what if the markets do worse or better than the median over the next 15 years? How do these possibilities affect Joe's retirement plans?

To help measure the impact of various possible outcomes, we introduce the notion of Joe's "relative standard of living" in retirement. We define this to be the ratio of Joe's net income after retirement to his net income in his last working year prior to retirement. Joe's goal is to have this ratio be at least 1.0. A ratio less than 1.0 is bad. A ratio greater than 1.0 is good.




Figure 2: Cumulative Density Function with Extra Monthly Savings

Before he retires, Joe's net income is $92.5\% - x$ measured as a percentage of his gross salary. After he retires, his net income is his social security income plus the 5% withdrawal from his retirement portfolio. Let p be the value of Joe's portfolio when he retires. Then his net income after retirement is $0.05p + 0.2$, again measured as a fraction of his gross salary. We divide to get:

$$\text{Relative standard of living} = \frac{0.05p + 0.2}{0.925 - x}$$

Consider the 10th and 90th percentiles in the cumulative density graph in Figure 2. The ending values p are 6.89 and 24.36 respectively. We have $x = 9.3\% = 0.093$. Plugging in these values in our equation gives relative standard of living ratios of 0.65 and 1.70.

Thus, if the markets do unusually poorly over the next 15 years (at the 10th percentile or worse), Joe must accept only 65% or less of his desired standard of living after retirement, assuming he is not willing to continue working after age 66. If Joe is concerned about this possibility, which he probably should be, he may want to save a little bit more.

On the other hand, if the markets do well over the next 15 years (at the 90th percentile or better), Joe can retire with 1.7 times his target standard of living, or even more. Joe is quite thrilled at this prospect!

In our analysis of Joe's retirement plan we see once again that uncertainty rules the world of investing. Joe is only 15 years away from retirement, yet we can only give him an 80% chance that his post-retirement standard of living will be somewhere between 65% and 170% of his pre-retirement standard of living. This is a wide range of possible outcomes, and it only covers 80% of the possibilities!

3.1 Social Security

Notice how important social security is in this analysis. Even though we are estimating Joe's social security benefits at only 20%, this seemingly small amount acts as a significant buffer which cushions the impact of the downside scenario. As an exercise, the reader may like to rework Joe's model under the assumption that there will be no social security at all when he retires. Not only does Joe have to increase his savings dramatically, the downside also becomes much more unpleasant under this assumption.

Another interesting and certainly relevant experiment would be to rework the model based on the assumption that the government invests all or part of Joe's future social security benefits in the stock market.

3.2 Rebalancing

Financial experts advise us to "rebalance" our portfolios occasionally. For example, Joe has an asset allocation of 5% cash, 25% bonds, and 70% stocks. As time goes on, because different segments of the market generate different returns, his portfolio will become out of balance with these percentages. Thus he should every once in a while transfer funds between his investments to bring his portfolio back into balance.

The model we developed for Joe assumes that he rebalances with exactly the same frequency as his periodic savings, which is monthly. Most people don't rebalance this often. Rebalancing once or twice per year is more common. The error caused by this discrepancy in our model is minor, however, and in any case it is overwhelmed by the inaccuracies of all of our other assumptions and estimates.

3.3 Salary Growth

As with rebalancing, our model assumes that Joe gets his raise in small installments every month. This is of course an unrealistic assumption. Once again, however, the error caused by this improper assumption is minor and we can safely ignore it.

3.4 Asset Allocation and Risk/Return Tradeoffs

It is interesting to experiment with our parameter estimation equations with different asset allocations. For example, the following table shows the estimated values of μ and σ for various mixtures of stocks and bonds. In this example we apply our estimation equations for nominal returns and no investment expenses.

	All bonds	75% bonds 25% stocks	50% bonds 50% stocks	25% bonds 75% stocks	All stocks
μ	4.7%	6.3%	7.7%	8.8%	9.7%
σ	7.9%	8.2%	11.0%	15.0%	19.5%

Notice that as the expected return μ increases, the standard deviation σ also increases. This is one of the fundamental laws of the marketplace: risk and return always go hand-in-hand.

In Joe's retirement plan, we could experiment with changing his asset allocation. More aggressive portfolios have greater median ending values, but they also have a wider spread of ending values (they are riskier). We leave these experiments as an exercise for the reader.

3.5 Investment Expenses

Recall that Joe's investment expenses were 35 basis points, and this was a parameter in our estimation of μ and σ . These are low expenses. The average expenses for the mutual fund industry are about 100 basis points higher, in the neighborhood of 1.35%. Let's suppose that Joe's expenses are more in line with the industry average at 135 basis points instead of only 35 basis points. What impact does this have on the model?

In our parameter estimation, the increase of 100 basis points in expenses brings our expected continuously compounded real rate of return after expenses relative to salary growth down from 3.7% to 2.8%, and it decreases the median ending value of Joe's portfolio after 15 years from 12.6 to 11.2. This is 1.4 years of Joe's gross salary! That's an enormous difference. Clearly expenses are tremendously important in our model.

John Bogle is the founder of the Vanguard Group of mutual fund investment companies and one of the founders of index funds. He is a champion of lower costs in the mutual fund industry. Bogle likes to say that "costs matter," and he's right.³

³See Bogle's book [4] and his collection of speeches [3].

3.6 Assumptions and Possible Directions for Future Work

We have made several strong simplifying assumptions about Joe's retirement plan.

We assumed that Joe does not vary his asset allocation over the 15 years remaining until his planned retirement. This assumption is called "constant relative risk aversion." As the value of Joe's portfolio grows and shrinks over time and he becomes wealthier or poorer, he keeps the proportion of his total portfolio value which is exposed to risky assets a constant.

This is a strong assumption. There is no reason to assume that all investors have constant relative risk aversion. Some investors may prefer to reduce their relative exposure to risk as they accumulate more wealth. These investors become increasingly concerned about preserving their accumulated capital as they get wealthier. Other investors may prefer to increase their relative risk exposure as they accumulate more wealth, perhaps based on the notion that at some point they have accumulated more than they need, and are hence willing to take on extra risk with the excess. There's no single strategy which can be defined as "correct" in the sense that it is the strategy all rational investors should follow.

Our random walk model can be modified to accommodate these more complicated kinds of relative risk aversion. Instead of being constants, our parameters μ and σ become functions of the portfolio value s .

We also assumed that Joe saves the same constant fraction of his gross salary every month. For some investors, as they get older and their children leave home and they pay off their mortgages (for example), they may find that they are able to save a much larger fraction of their gross salary for retirement. To accommodate these kinds of situations, in our model we must change the parameter k to be a function of time instead of a constant.

We have also assumed that Joe's future salary growth rate is some constant amount in excess of inflation. This assumption does not take into account the variability of this excess amount or allow for the modeling of salary growth rates which we expect to increase or decrease as Joe gets older. It would be possible to enhance the model for these factors by making the salary growth rate in excess of inflation a random variable with its own estimated growth rate and volatility parameters.

It would perhaps be an interesting future project to make these changes to the computer program we use to do Monte Carlo simulations of the model. For example, given an arbitrary "utility function" U , we could modify the program so that at each step in the simulation we recalculate μ and σ so as to maximize the expected utility of the portfolio ending values at the target retirement date.⁴ The algorithms required to implement this idea in a computer program in such a way as to still be able to do a large number of Monte Carlo simulations in a

⁴For a discussion of "utility functions," see [5].

reasonable amount of time are a bit tricky, but it's doable.

On an even more ambitious scale, it is certainly conceivable that we could enhance the model and the program to do a complete life-cycle plan. Such a model and simulation might, for example, combine the pre-retirement and post-retirement analysis, take into account Joe's human capital during his working years as part of the risk aversion machinery, and accommodate his desire (if any) to leave a bequest for his heirs.

4 Post-Retirement Planning

In our retirement model for Joe we made the unsupported assumption that 5% was a reasonable inflation-adjusted portfolio withdrawal rate in retirement. In this section we fulfill our promise to return to address this issue.

The greatest risk faced by retirees who are living off an investment portfolio is outliving their money. We can use the model we developed for pre-retirement planning to analyze this post-retirement problem. Prior to retirement, we set our new parameter k to a positive constant number to model periodic additions to the portfolio. After retirement, when we are making periodic withdrawals instead of additions, we simply set k to a negative constant.

Let's use our model to test our 5% inflation-adjusted withdrawal rate and see what happens.

First we need to set the parameters for the model. Let's assume the same portfolio we used in the previous section: 5% cash, 25% bonds, and 70% stocks, with the same expense ratio of 35 basis points. Instead of estimating real returns relative to Joe's salary growth, however, we will simply estimate real returns relative to inflation. We apply our parameter estimation equations to get the following values for μ and σ :

$$\begin{aligned}\mu &= 5.1860\% \\ \sigma &= 14.7012\%\end{aligned}$$

We will use 30 years for the t parameter. This may seem large. Joe is planning to retire at age 66, so we're assuming he will live to age 96. This is certainly considerably greater than Joe's life expectancy. Assuming that Joe is in good health when he retires, however, living to age 96 is not at all unlikely, and because running out of money while Joe is still alive is such a serious problem, we are planning conservatively.

We're interested in the probability that Joe's retirement portfolio will last for 30 years given inflation-adjusted withdrawals of 5% per year. For this purpose, we can use any starting value s_0 for the portfolio as long as we also set the parameter k to 5% of s_0 . We'll use $s_0 = 100$ and $k = -5$.

We'll assume that Joe takes out living expenses monthly from his portfolio, so we set the parameter $\Delta t = 1/12$.

We now have all our parameters and can run our simulation to graph the cumulative density function for $s(30)$. The result is shown in Figure 3.

As you can see, about 27% of the ending values are less than 0, which means that in 27% of the 100,000 simulations Joe's portfolio ran out of money before 30 years. This is risky indeed, and it's a serious problem.

Note that it's not all bad news. There's an upside to this story. The median ending value is about \$100, which is the same as the starting value. At the 90th




Figure 3: Post-Retirement Cumulative Density Function

percentile, the ending value is about \$600, which means that Joe's portfolio has actually grown by a factor of 6 over 30 years, even after adjusting for inflation, even though he's taking money out every month!

Again, uncertainty rules, even more dramatically than in our pre-retirement planning.

Despite the upside, Joe is concerned about the downside, as he should be. He's interested in what would happen with other possible withdrawal rates. Perhaps a somewhat smaller withdrawal rate would bring the failure rate down to a more acceptable level.

We can repeat our simulation with other withdrawal rates. Figure 4 shows the result, with withdrawal rates ranging from 3% to 9% in steps of 0.5%.

Clearly there is risk even with a low withdrawal rate of 3%, where the failure rate is 4%. But this is certainly better than the 27% failure rate with a withdrawal rate of 5%.

Why doesn't Joe simply plan for a withdrawal rate of 3%? Why did we choose the riskier 5% in his retirement plan?

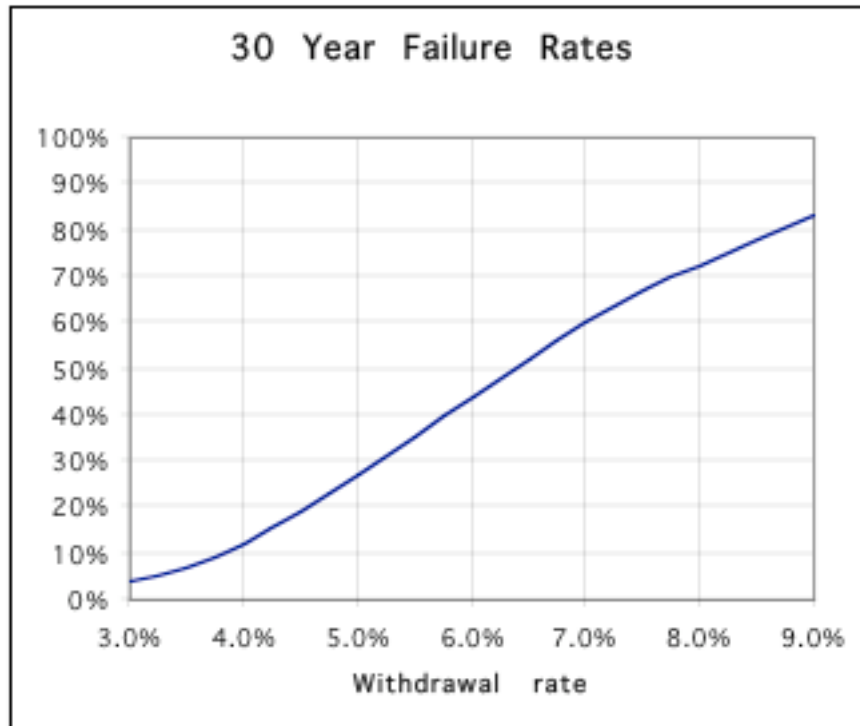


Figure 4: 30 Year Failure Rates

Let's see what would happen with a withdrawal rate of 3% instead of 5%. Recall that in Joe's retirement plan we had to multiply the amount of salary he needed to replace via his savings by 20 to account for the 5% withdrawal rate. The multiplier 20 is simply $1/5\%$. If we use a 3% withdrawal rate, we must instead multiply by $1/3\% = 33$. This is a much bigger multiplier than 20. We invite the reader to rework Joe's retirement plan under this assumption. You will quickly discover that Joe must save a ridiculously large fraction of his gross income to reach the new goal.

Thus Joe can protect himself against poor market conditions after he retires with this plan only at the cost of making enormous financial sacrifices while he is still working. If the markets fail to do poorly, Joe ends up dying enormously wealthy without ever having had a chance to enjoy his money! Economists call this "an inefficient tradeoff of present versus future consumption."

Joe has to make reasonable compromises. It's not at all an easy problem or an easy decision.

There are several things Joe could do to try to deal with this problem.

First, Joe could adopt a different retirement strategy. Instead of withdrawing a constant inflation-adjusted amount every year from his portfolio, he could perhaps try to withdraw less than usual during years in which the market isn't doing well. There are a number of possible algorithms that could implement this strategy. All of them require Joe to be willing to plan for income levels which fluctuate over a wide range in retirement. It is not possible to fund a steady relatively stable income stream over a long period of time with any of these strategies.

A second thing Joe could do is use all or part of his retirement savings when he retires to purchase a lifetime annuity from an insurance company. This is in fact a popular decision made by many retirees. The insurance company guarantees a steady income stream for as long as Joe lives. Some kinds of annuities even offer provisions to keep up with inflation. If Joe plans to do this, we must find out how much it will cost to provide each dollar of guaranteed income and factor that into the model. This number is called the "annuity factor." Incorporating this into the model may increase or decrease Joe's target savings amount, depending on a variety of factors.

The advantage of an annuity is that it provides protection against outliving your money in retirement. The disadvantages are less flexibility and a smaller estate or no estate at all to pass on to your heirs.

One strategy might be to purchase an annuity that is large enough so that the annuity income together with social security income provides a minimum floor level of income to meet basic living needs.

These issues are complicated indeed. Our models and our software help to investigate them, but they do not provide a definitive solution to the problem. Indeed, it appears that the problem has no completely satisfactory solution. Once again, the fickleness of the markets has conspired to make planning extraordinarily difficult and uncertain.⁵

⁵For another perspective on this problem see [1].

References

- [1] William J. Bernstein. The retirement calculator from hell.
<http://www.efficientfrontier.com/ef/998/hell.htm>.
- [2] Zvi Bodie and Robert C. Merton. *Finance*. Prentice-Hall, preliminary edition, 1998.
- [3] John C. Bogle. Speeches.
http://www.vanguard.com/bogle_site/bogle-speeches.html.
- [4] John C. Bogle. *Bogle on Mutual Funds: New Perspectives for the Intelligent Investor*. Dell Publishing, 1994.
- [5] John Norstad. An introduction to utility theory.
<http://www.norstad.org/finance>, Mar 1999.
- [6] John Norstad. Random walker.
<http://www.norstad.org/finance>, Feb 2005.
- [7] John Norstad. Random walks.
<http://www.norstad.org/finance>, Jan 2005.