Will It Last?

How to estimate sustainable spending rates using a handful of assumptions and a Microsoft Excel function.

QUANT U

Paul D. Kaplan

What do a retiree and charitable foundations have in common? They both need to decide how much to spend each year over their lifetimes so as not to go broke in the process. In this, the main difference between them is time horizon. For the retiree, the time horizon is until an unknown date of death. For the foundation, the time horizon is infinite, because foundations are typically set up to be run forever.

The problem of finding a sustainable spending rate, at least for a finite horizon, has been solved in a few ways. Some approaches use Monte Carlo simulation. But another approach, developed by Milevsky and Robinson (2005) and, in more detail, Milevsky (2006, chapter 9), does not use Monte Carlo simulation. Rather, by making some simplifying assumptions, they show how to calculate the probability of running out of money too soon from a handful of parameters and a little known function in Microsoft Excel.

In this edition of Quant U, I first present the Milvesky-Robinson model. Then I show step by step how to develop the assumptions and apply them to the model. In doing so, I will draw on some of the formulas that I presented in previous issues of Quant U. Finally, I show how the model can be used as an aid in making asset-allocation decisions.

Stochastic Present Value

Milvesky and Robinson start with the concept of present value. The present value of $1 over a period of future years tells us how much money initially a portfolio must have to allow $1 to be withdrawn each year, assuming that each year the portfolio earns the discount rate. If the discount rate is constant, we have:

1 $$PV = \sum_{t=1}^{T} (1+R)^{-t} = \frac{1-(1+R)^{-T}}{R}$$

Where PV is the present value of $1 over a period of T years and R is the discount rate or rate of return. In the case of a foundation, T becomes ∞ so that $PV=1/R$. So, if $R=2\%$, we would need 1/0.02, which is $50 in the initial portfolio for each $1 to be spent each year.

Of course, if the portfolio is invested in risky assets, the rate of return each year is unknown ahead of time. Also, in the case of the retiree, the final year is unknown. Substituting the known value of R and T for unknown returns ($\tilde{R}_1, \tilde{R}_2, \ldots, \tilde{R}_{\tilde{T}}$) and, in the case of the retiree, an unknown final year ($\tilde{T}$), Milveksky and Robinson introduce the concept of *stochastic present value*:

2 $$\widetilde{SPV} = \sum_{t=1}^{\tilde{T}} \prod_{v=1}^{t} (1+\tilde{R}_v)^{-1}$$

What makes $\widetilde{SPV}$ useful is that under certain assumptions about the distribution of returns and the distribution of the final year, it has a specific distribution, namely that of a *reciprocal gamma distribution*.

A reciprocal gamma distribution takes two parameters, α and β. The probability density function is:

3 $$f_{\widetilde{SPV}}(x) = \frac{x^{-(\alpha+1)} exp(-\frac{1}{\beta x})}{\Gamma(\alpha)\beta^{\alpha}}$$

where $\Gamma(.)$ is the gamma function. If in Excel you have a cell named *alpha*, the formula =EXP(GAMMALN(alpha)) gives you the value of the gamma function with *alpha* as the argument. (This is not the function I referred to earlier, although it is related.)

Below, I explain how to derive the values of α and β from the assumptions regarding returns and the final year. But before doing that, I think that it is instructive to look at the probability density curves of some reciprocal gamma distributions, which represent stochastic present value for some specific cases. In **EXHIBIT 1**, I show these curves in two cases: (1) a 65-year female retiree and (2) a foundation.

In both cases, I assume that the investment portfolio is 50% stocks/50% bonds. The only difference is in the assumption regarding the final year. In the case of the 65-year-old female retiree, the values for α and β parameters reflect mortality assumptions for a 65-year-old woman. In the case of the foundation, there is no final year, as the period is infinitely long. This should mean that the foundation needs considerably more funds in its initial portfolio per $1 spent annually than does the retiree. We can see that this is indeed the case as the probability density curve is fairly far to the right than that of the retiree.

EXHIBIT 1

Probability Distribution Functions for Stochastic Present Value

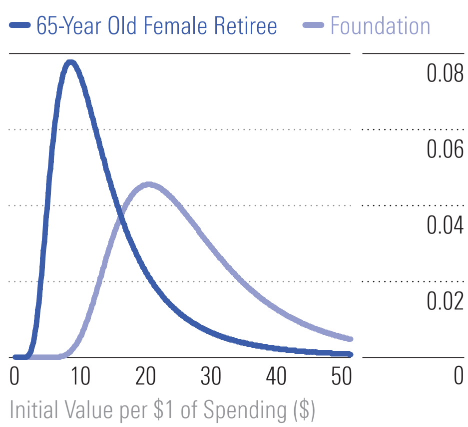

Source: Morningstar.

EXHIBIT 2

The Gompertz and Exponential Models of Mortality

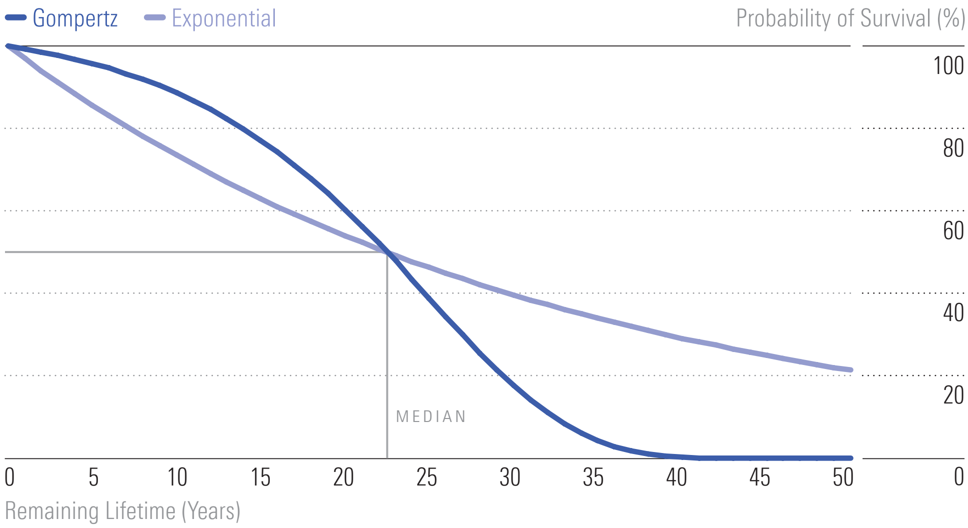

Source: Morningstar.

The Mortality Assumption: Exponentially Distributed Remaining Life

The Milevsky-Robinson model assumes that the remaining lifetime of a person follows an exponential distribution. This means that the probability of the remaining lifetime being at least t is:

❹ $$Pr[\tilde{T} \geq t] = exp(-\lambda t)$$

where λ is a parameter to be determined. This is very different than the Gompertz model that Milevsky (2006, chapter 3, and 2012) presents and that I discuss in Quant U in the February/March 2015 issue.

In **EXHIBIT 2**, I plot the probabilities of remaining lifetime for both models. For the Gompertz model, I use the parameters for a 65-year-old woman, which appear in Blanchett and Kaplan (2013). For the exponential model, set λ such that the median remaining lifetime under the exponential model is the same as under the Gompertz model. (I explain why and how to do this below.)

As you can see, the models are quite different. The backward S-shaped curve produced by the Gompertz model provides fairly realistic probabilities of survival over a wide range of possible remaining lifetimes. In contrast, the convex curve produced by the exponential model, at the median, are not realistic.

So, why use the exponential model? Because the mathematics of the Milevsky-Robinson model require it. Fortunately, it works well in the context of that model. As Milevshy and Robinson explain:

"Although human aging does not conform to an exponential or constant force of mortality assumption—which means that death would occur at a constant rate—for the purposes of estimating a sustainable spending rate, it does a remarkably good job when properly calibrated."

As they say, the way to properly calibrate it is to set λ such that the median remaining lifetime matches that of a realistic value. To do this, I suggest starting with the Gompertz model, which can be stated as:

❺ $$q = exp\left\{\left(1-exp\left(\tfrac{t}{b}\right)\right)exp\left(\tfrac{a-m}{b}\right)\right\}$$

where

q = the probability of surviving at least t more years

a = current age in years

m = the mode of the distribution of the age of death

b = The dispersion of the age of death around the mode

The Gompertz formula can also be written to find t given q:

❻ $$t = b \cdot ln\left\{1-\frac{ln(q)}{exp\left(\frac{a-m}{b}\right)}\right\}$$

To find the median remaining lifetime, set q=0.5. Because we are modeling a 65-year-old woman, we set a=65 and use the values of m and b estimated by Blanchett and Kaplan (2013), namely:

m = 90

b = 8.63

Plugging these values into equation 6 yields the median remaining life:

t_{Med} = 22.50

To find λ, we make t_{Med} the median remaining life of the exponential model:

❼ $$0.5 = exp(-\lambda t_{Med})$$

Solving for λ, we have:

❽ $$\lambda = \frac{ln(2)}{t_{Med}} = \frac{0.6931}{t_{Med}}$$

Plugging our value for t_{Med} into equation (8), we have:

$\lambda = 0.0308$

In the case of the foundation, none of this is necessary. We simply set λ=0.

The Investment Return Assumption: The Lognormal Distribution

In Quant U in the August/September 2015 issue, I explained why the lognormal distribution is the natural distribution for modeling returns (assuming that return distributions do not have fat tails). We say that a random variable representing returns, $\tilde{R}$, follows a lognormal distribution if $ln(1+\tilde{R})$ follows a normal distribution. Denote the expected value of $ln(1+\tilde{R})$ as μ and the standard deviation of $ln(1+\tilde{R})$ as σ. These lognormal parameters, μ and σ, along with the exponential mortality parameter, λ, are the inputs to the Milevsky-Robinson model.[1]

Given the expected value of $\tilde{R}$, M, and the standard deviation of $\tilde{R}$, S, we can find μ and σ as follows:

(9) $$\sigma = \sqrt{ln\left(1+\left(\frac{S}{1+M}\right)^2\right)}$$

(10) $$\mu = ln(1+M) - \frac{1}{2}\sigma^2$$

To illustrate how to use these formulas to form the inputs to the Milevsky-Robinson model, I use a very simple asset-allocation model in which there are just two asset classes, equity and fixed income.

Let

M_E = expected return on equity
S_E = standard deviation of equity return
M_F = expected return on fixed income
S_F = standard deviation of fixed-income return
ρ = Correlation between equity and fixed-income returns
θ = Fraction of the portfolio in equity

The expected return and standard deviation of return on a portfolio are given by:

(11) $$M = \theta M_E + (1-\theta)M_F$$

(12) $$S = \sqrt{\theta^2 S_E^2 + (1-\theta)^2 S_F^2 + 2\theta(1-\theta)S_E S_F \rho}$$

For example, suppose that the parameters of the asset-allocation model have the following values:

M_E = 7%
S_E = 20%
M_F = 3%
S_F = 10%
ρ = 20%
θ = 50%

Plugging these values into equations 11 and 12, we have:

M = 5.00%
S = 12.04%

EXHIBIT 3

The Trade-Offs Between Spending Rate and Probability of Success

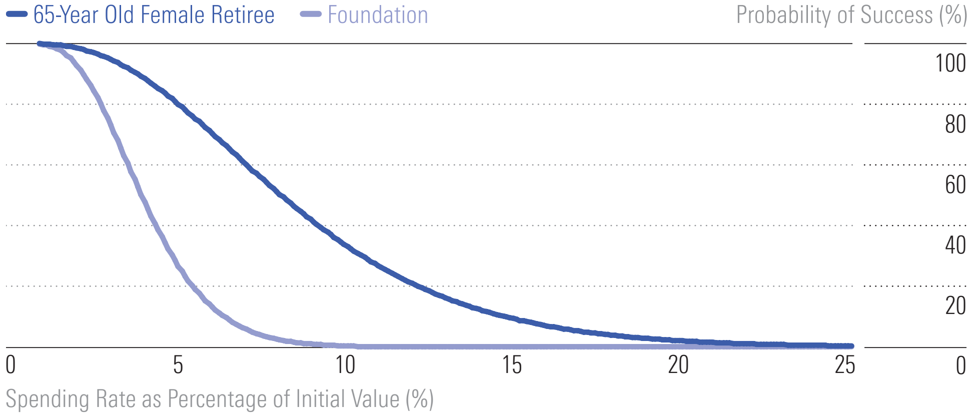

Source: Morningstar.

Plugging these values into equations 9 and 10, we obtain the values of the parameters of the lognormal distribution that we need for the Milevsky-Robinson model:

μ = 4.23%
σ = 11.43%

Calculating the Probability of Success

Recall that stochastic present value, $\widetilde{SPV}$, under the assumptions that I have presented has a reciprocal gamma distribution. This means that its reciprocal follows a gamma distribution. Because Excel contains a function that returns values of the cumulative probability distribution function of gamma distributions, implementing the Milevsky-Robinson model in Excel is quite practical.

Let s denote the amount of annual spending expressed as a fraction of initial wealth. For example, if spending is $40,000 per year and initial wealth is $1 million, s=4%. To spend at this rate without going broke, we need for $\widetilde{SPV} \le \frac{1}{s}$. This condition is equivalent to the condition that $\frac{1}{\widetilde{SPV}} \ge s$. Because $\frac{1}{\widetilde{SPV}}$ follows a gamma distribution, we have:

(13) $$Pr\left[\frac{1}{\widetilde{SPV}} \ge s\right] = 1 - F_{\Gamma}(s; \alpha, \beta)$$

Where $F_{\Gamma}(.;\alpha,\beta)$ is the cumulative distribution function for a gamma distribution with parameters α and β. I refer to the probability of not running out of money as the *probability of success*.[2]

The values of α and β come from the values of the mortality parameter (λ) and the lognormal parameters of investment returns (μ and σ) as follows:

(14) $$\alpha = \frac{2\mu + \sigma^2 + 4\lambda}{\sigma^2 + \lambda}$$

(15) $$\beta = \frac{\sigma^2 + \lambda}{2}$$

Using the values that I came up with earlier, for the 65-year-old female retiree:

$$\alpha_R = \frac{2 \cdot 0.0423 + 0.1143^2 + 4 \cdot 0.0308}{0.1143^2 + 0.0308} = 4.0328$$

$$\beta_R = \frac{0.1143^2 + 0.0308}{2} = 0.0219$$

For the foundation, λ=0, so that:

$$\alpha_0 = \frac{2 \cdot 0.0423 + 0.1143^2}{0.1143^2} = 6.4682$$

$$\beta_0 = \frac{0.1143^2}{2} = 0.0065$$

1 In Milevsky and Robinson (2005) and Milevsky (2006), μ denotes $lnE[1+\tilde{R}]$ whereas I use this symbol to denote $E[ln(1+\tilde{R})]$. The difference between these two values is $\frac{1}{2}\sigma^2$. So, when using their formulas, I replace their μ with my $\mu + \frac{1}{2}\sigma^2$.

2 Milevsky and Robinson calculate the probability of ruin, which is the probability of going broke. The probability of ruin is simply 100% less the probability of success.

EXHIBIT 4

Trade-Off Curves for the Retiree at Various Allocations to Equity

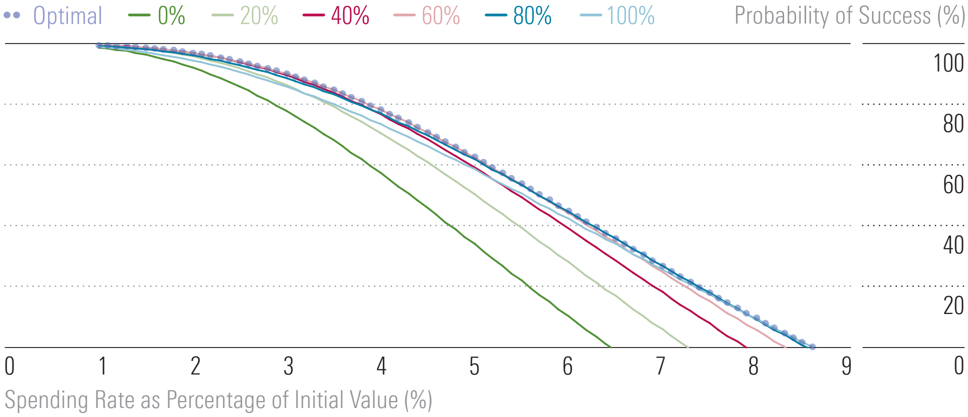

Source: Morningstar.

EXHIBIT 5

Optimal Equity Allocations for the Retiree

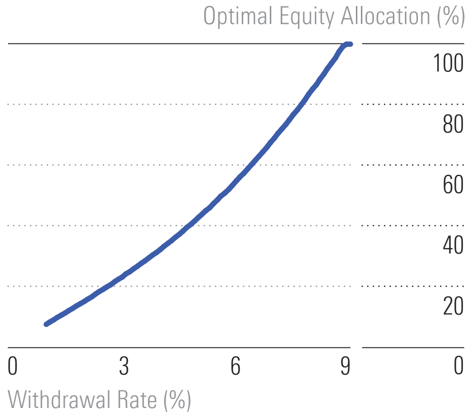

Source: Morningstar.

Now, we just need to calculate the value of the cumulative gamma distribution function with these parameters. In Excel, =GAMMADIST(s,alpha,beta,TRUE) returns the value of the cumulative probability function of a gamma distribution at *s* with parameter values alpha and beta. Hence, to calculate the probability of success at spending rate *s*, use =1-GAMMADIST(s,alpha,beta,TRUE). Using the values that I present above, for the retiree, to find the probability of not running out of money with a 4% spending rate, type =1-GAMMADIST(0.04,4.0328,0.0219,TRUE) into a cell. The result is 89.1%. For the foundation, use =1-GAMMADIST(0.04,6.4682,0.0065,TRUE). The result is 49.7%.

There is a trade-off between the spending rate and the probability of success: the greater the spending rate, the less likely that the portfolio will last long enough. To illustrate, I plot in **EXHIBIT 3** various spending rates versus their probability of sucess for the retiree and foundation. Because of its infinite time horizon, the foundation's curve is far to the left of the curve for the retiree.

Asset Allocation and the Probability of Success

So far, I have taken the investment decision, which in the model I present here is allocation to equity (θ), as given. However, a spending strategy is really a joint decision of how much to spend each year and how to invest the portfolio. In **EXHIBIT 4**, I draw the trade-off curves between the spending rate and the probability of success for the retiree using various levels of equity allocation; namely 0%, 20%, 40%, 60%, 80%, and 100%. For each spending rate, we can see which level of equity allocation results in the highest probability of success. While equity allocations of 0% and 20% are never the best, allocations of 40%, 60%, 80%, and 100% are the best at different levels of spending, with the best allocation increasing with the spending rate.

I took this analysis one step further by finding for each spending rate from 1% to 9% in increments of 0.1%, the equity allocation that maximizes the probability of success. Given the large number of optimizations, for speed I did the calculations outside of Excel.[3] I present the results in **EXHIBIT 4** and **EXHIBIT 5**. In **EXHIBIT 4**, the dotted curve shows the maximum probability of success for each spending rate. **EXHIBIT 5** shows the optimal allocation to equity at each spending rate. At the low end, it is 44.5% for a spending rate of 1%. It steadily increases until it reaches 100% for a spending rate of about 9%.

Because the optimal allocation to equity increases with the spending rate, the more aggressive the rate of spending, the more aggressive the optimal portfolio. This is not to say that one should invest retirement savings aggressively, because the probability of success declines with spending rate, even at the optimal level of equity allocation.

Simplified in Excel

The main difference in calculating a sustainable spending rate for institutions and individuals is time horizon. A person's time horizon is unknown; an institution's is infinite. One approach to finding the sustainable rate is to use Monte Carlo simulation. As I've shown, another way is to make some simplifying assumptions and solve it analytically in Microsoft Excel. M

Paul D. Kaplan, Ph.D., CFA, is director of research with Morningstar Canada. He is a member of the editorial board of *Morningstar* magazine.

References

Blanchett, David and Paul Kaplan. 2013. "Alpha, Beta, and Now… Gamma." *Journal of Retirement*. Fall.

Milevsky, Moshe A. 2006. *The Calculus of Retirement Income*. New York, N.Y.: Cambridge University Press.

Milevsky, Moshe A. 2012. *The 7 Most Important Equations for Your Retirement*. John Wiley & Sons Canada, Ltd.

Milevsky, Moshe A. and Chris Robinson. 2005. "A Sustainable Spending Rate without Simulation." *Financial Analysts Journal*, November/December.

3 I used the fmincon function in MatLab. This could have been done in Excel using the Solver and writing a routine in VBA to loop over the values of s and save the result, but I found it easier to do in MatLab.