FINANCIAL
SERVICES
REVIEW

Financial Services Review 14 (2005) 73–84

Long-range confidence interval projections and probability estimates

William J. Trainor, Jr.*

Department of Finance, Western Kentucky University, Bowling Green KY 42101, USA

Abstract

Accurately estimating long-range confidence intervals and probability estimates associated with an investment portfolio is critical to the financial planning process. To address this, computationally intensive data limiting Monte Carlo techniques have become popular. The results of this study suggest that simple theoretical probabilistic return projections based only on an expected return and standard deviation estimate are just as accurate. Thus, using a simple spreadsheet, financial planners and investors can accurately assess a wide range of possible outcomes for a myriad of asset allocations, investment amounts, and time horizons without needing to resort to more sophisticated methods.

JEL classifications: C15

Keywords: Forecasting; Confidence intervals; Bootstrapping

1. Introduction

Accurately projecting expected outcomes and the risk associated with achieving those outcomes is the primary step to creating a financial plan that hopefully will meet an investor's long-term needs. Being able to determine the probability of reaching a particular goal is critical to the investment process. The minimum probability of success often dictates the amount of contributions needed, the asset mix, and the expected time to retirement. Thus, projecting future wealth values and the probability associated with reaching them is critical to successful long-term planning. An investor whose goal is $500,000 for retirement in 20

* Tel.: +1-270-745-6337.
E-mail address: Bill.Trainor@WKU.EDU

years needs to know not only what return is required to reach that goal, but also the probability that the goal may not be reached. The second piece of information is just as important as the first in determining both the asset mix and the magnitude of the contributions needed to reduce the risk of a shortfall.

Currently, there are two major avenues of approach when projecting long-term retirement values and assessing their probability of success. The first involves estimating an expected return and standard deviation, projecting a terminal wealth value, and computing a confidence interval around that terminal value based on the standard deviation estimate. Concurrently, the probability of reaching a particular value can also be calculated. The second procedure is value at risk, or VAR analysis. VAR analysis simply rearranges the standard deviation to define the maximum loss or minimum gain that can be realized given a fixed confidence level. In essence, it gives a dollar value that is at risk given some particular probability level.

In applying either of the above approaches, it is generally assumed that asset returns are normally distributed. This assumption is made for tractability although it is well known that equity returns demonstrate some skewness and much more kurtosis (extreme returns) then would generally be expected from a normal distribution (Mandlebrot, 1963; Affleck-Graves & McDonald, 1989; Zhou, 1993). Many researchers have considered this problem to be severe enough that typical beta estimates used to test finance's benchmark model, the Capital Asset Pricing Model, are thought to be misspecified (Knez & Ready, 1997; Shalit & Yitzhaki, 2002), and some have even suggested alternative distributions need to be assumed (Hodgson, Linton, & Vorkink, 2002).

On the financial planning front, rather than assuming normality of returns, the use of computerized simulation methods has become popular. The most common method applied is actually called bootstrapping which can be thought of as a type of non-parametric Monte Carlo analysis.[1] By running thousands of simulations that take the underlying data as given, bootstrapping methods can explore thousands of possible return paths, and derive confidence intervals encompassing these return paths. No assumption about the underlying distribution or its properties is assumed. However, bootstrapping does make the assumption that future paths will have the same basic historical return realizations that have been experienced in the past. Thus, if the historical data does not have an 80% monthly return, bootstrapping methods will not create one.

The use of bootstrapping techniques has been used extensively by researchers for a host of applications. Booth (2004) used bootstrapping to study retirement targets, Hickman et al. (2001) and Mukherji (2003) employed it to determine optimal retirement allocations, while Pye (2000) and Milevsky (2001) used bootstrapping to determine optimal withdrawal rates at retirement.

Although the use of bootstrapping allows researchers to eliminate assumptions about the return distribution, it is not always the most efficient procedure given conditional accuracy requirements. In addition, bootstrapping is computationally intensive, the results are subject to the time period studied or for which data are available, and there is the implicit assumption that returns outside the historical record are not possible.[2] Thus, applying bootstrapping methods in the name of distributional robustness relative to theoretical models may not always be the most efficient or effective approach.

In fact, the results of this study show that for a simple "vanilla" terminal wealth projection with associated confidence intervals, theoretically derived confidence intervals based on the assumption of normality coincide quite closely with computationally intensive bootstrapping methods. These results hold throughout the entire spectrum of simple equity/bond portfolios. Based on the fact that portfolios with additional asset classes are likely to have returns that appear even more normally distributed, these results would extend to more expansive portfolios as well.

Thus, rather than using bootstrapping methods for creating extended probabilistic terminal wealth projections, almost the exact same results can be attained from using a theoretical model. The benefits of this allow quick readjustment of parameters for "what if" scenarios that may deviate significantly from historical realizations. Also, results can be instantaneously derived for any changes in expected return/standard deviation assumptions without resorting to additional simulation runs. Bootstrapping methods do not have this flexibility.

The immediate consequence of these results suggest that investors and financial planners need only arm themselves with an expected return and standard deviation estimate for their portfolios to accurately forecast expected returns and the associated probability estimates that surround that expectation. In addition, probability forecasts associated with portfolios composed of multiple asset classes can be easily analyzed which may not be the case when using bootstrapping. A portfolio structure such as 20% international, 50% U.S. equity, 20% high yield bonds, and 10% REITs would be difficult to analyze using bootstrapping methods since historical data for some of these asset classes is limited. However, deriving an expected return and a standard deviation for this portfolio is a relatively more simplistic exercise.

In fact, using Morningstar's Principia, an investor or financial planner can actually input the actual mutual funds held or that are under consideration and attain return and standard deviation estimates for the combined holdings. Regardless of how one derives the two estimates, once so armed, long-range return projections and the associated probability estimates are easy to calculate and should be just as accurate as bootstrapping procedures.

In defense of bootstrapping methods, it should be noted that applying theoretical models relying on the normality assumption will certainly not be optimal in all cases. For more complicated portfolio structures, and certainly those that may exhibit extreme non-normality, bootstrapping methods are likely to be preferred. Portfolios that are highly leveraged or have large short positions may fall into this category (Kritzman & Chow, 2001). However, for most investors or financial planners, this should not be a concern.

2. Forecasting approximate long-range confidence intervals

An initial forecast usually starts by estimating expected returns and standard deviations that are generally annualized in one fashion or another. This is usually performed by annualizing daily, weekly, or monthly data. When annualizing data, the nominal or effective rate can be calculated by multiplying or taking the exponential of the return by the number of periods in a year. For annualizing the standard deviation, the common rule of thumb is to multiply the standard deviation by the square root of the number of periods in a year. This is only an approximation because of the fact that the nature of compounding causes the dispersion of returns to follow a lognormal distribution as opposed to a normal distribution.

As an example, suppose the average monthly return is found to be 1% while the monthly standard deviation is 4%. The nominal annual return is easily calculated as 12% ($12 \times 1\%$) or the effective rate can be calculated as $1.01^{12} - 1 = 12.68\%$. The annualized standard deviation is then computed as $4\% \times 12^{0.5} = 13.86\%$.

Whichever method is used to arrive at the annual return and standard deviation numbers, calculating a confidence interval for single period forecasts under the assumption of normality is a simple procedure of adding and subtracting a certain number of standard deviations to the expected mean where the number is determined by the magnitude of the confidence level. The common confidence level is 95% which is two standard deviations (theoretically 1.96 is the exact number for large samples) added and subtracted to the mean.

However, this procedure is only correct for single period forecasts. When extending these forecasts, the common rule of thumb is to divide the standard deviation by the square root of the number of periods into which the forecast is being made, and again adding and subtracting this adjusted standard deviation to the mean. This is the same procedure used to annualize the standard deviation above, except that we now divide by the square root as opposed to multiplying since the time frame is being extended beyond a year.

Extending the example above, assume that the annualized standard deviation is wanted for a 10 year forecast. The expected annualized mean remains 12.68%, but the annualized standard deviation around the mean is $13.86\%/10^{0.5} = 4.38\%$. Thus, the 95% lower and upper annualized limits for a 10 year horizon are $12.68\% \pm 1.96\ (4.38\%)$ or between 4.09% and 21.27%. Again this procedure is only an approximation since the return distribution becomes lognormal because of the compounding of returns. Note also that the annualized standard deviation decreases as the time horizon increases. This, of course, led to the controversy on the time diversification issue which will not be discussed here [see Samuleson (1963); Bodie et al. (2002)].

The approximation used above is remarkably accurate for limited time horizons under reasonable values for expected annual equity returns and standard deviations. As an example, using a 95% confidence interval, and assuming a mean of 11.66% and a standard deviation of 21.23% that correspond to the actual mean and standard deviation of equity returns from January 1926 to December 2002, the approximation procedure has less than a 15% error through 10 years. However, as Table 1 shows, this error increases to more than 35% by 20 years and to over 60% by 30 years. Thus, for long-term forecasting, it is not a recommended procedure.

One should also note that the lower wealth limits for both the theoretical and approximation methods initially decrease before monotonically increasing around year four. This is a direct result of the competing forces between risk and return. Initially, the higher standard deviation outweighs the positive expected return, but eventually this trend is reversed. The point of this reversal depends on the relative values of the expected return and standard deviation. As long as the expected return is positive, the probability of any loss approaches zero with an increasing time horizon regardless of the magnitude of the standard deviation. Table 1 demonstrates this by year 20 where the lower limit exceeds the initial investment.

Table 1
95% confidence intervals from 1 to 40 years for both approximate and theoretically correct procedures assuming an annualized mean and standard deviation of 11.66% and 21.23%, respectively. Lower and upper limits are shown on a per dollar basis with percentage differences between theoretical and approximation method in parenthesis.

Time	Lower limit-annual return	Upper limit-annual return	Lower limit per $	Upper limit per $
Theoretical				
1	−24.18%	58.71%	$.76	$1.59
5	−7.01%	29.40%	$.70	$3.63
10	−2.40%	23.29%	$.78	$8.11
20	1.00%	19.14%	$1.22	$33.20
30	2.54%	17.35%	$2.12	$121.38
40	3.47%	16.29%	$3.92	$418.74
Approximation			(% difference)	(% difference)
1	−29.95%	53.27%	$.70 (−7.61%)	$1.53 (−3.42%)
5	−6.95%	30.27%	$.70 (.31%)	$3.75 (3.42%)
10	−1.50%	24.82%	$.86 (9.60%)	$9.18 (13.15%)
20	2.36%	20.96%	$1.59 (30.58%)	$44.99 (35.54%)
30	4.06%	19.26%	$3.30 (55.55%)	$197.02 (62.31%)
40	5.08%	18.24%	$7.26 (85.35%)	$813.67 (94.31%)

This is also why Las Vegas casinos rarely lose money over any marginal period of time because from the casino's point of view, all their games have positive expected returns.

3. Calculating theoretically correct long-range confidence intervals

Although there is some loss of parsimony in calculating theoretically correct confidence intervals or probability estimates, it behooves the financial planner to do so when making projections much beyond 10 years. The process involves converting estimated discrete returns into their continuous counterparts, determining continuous confidence intervals, and then converting these values back into discrete values.

Specifically, both the expected return and standard deviation must be converted into their continuous counterparts. The equations to do this are the following:

$$\mu_c = \ln(1 + \mu_d) - \sigma^2{}_c/2 \tag{1}$$

$$\sigma_c = \sqrt{\ln[\sigma^2{}_d/(1 + \mu_d)^2 + 1]} \tag{2}$$

where μ_c = continuous expected return
μ_d = discrete expected return
σ_c = continuous standard deviation
$\sigma^2{}_d$ = discrete variance.

Note that the continuous mean is a function of the variance. This is because the lognormal distribution is not symmetric and the value of the variance affects the spread of the distribution that in turn affects the mean. Once the above values are found, the standard deviation associated with any time frame can be found by applying the following:

$$\sigma_{TAV} = \sigma_c / n^{.5} \quad (3)$$

where σ_{TAV} = the time adjusted standard deviation and n = the number of years into the future for which the confidence interval or probability estimate is being calculated. To find the range of continuous values (CV), just add and subtract the number of standard deviations associated with a particular confidence level:

$$CV = \mu_c \pm Z(\sigma_{TAV}) \quad (4)$$

where μ_c is the continuous expected return as defined in Eq. (1) and Z is found from the standard normal distribution table. For a 95% confidence level, Z would be 1.96 assuming an adequate number of observations.

To find the discrete confidence interval range, simply convert the continuous values back into discrete values (DV) by using the inverse natural logarithm function where e = 2.71828 and raise to the nth power:

$$DV = (e^{CV})^n \quad (5)$$

As an example, assume a mean of 10%, a standard deviation of 15% and let n = 10. By applying Eqs. (1) and (2), the continuous mean and standard deviation can be calculated as: $\mu_c = \ln(1 + .1) - 0.01842/2 = 0.08610$; $\sigma_c = \sqrt{\ln[.15^2/(1 + .1)^2 + 1]} = 0.13574$. Plugging σ_c into Eq. (3), $\sigma_{TAV} = 0.13574/10^{.5} = 0.04292$. To find the 95% range of returns, use Eq. (4) to find the continuous confidence interval that is 0.08610 ± 1.96 (0.04292) which then gives a return range between 0.00198 to 0.17022. Now convert these continuous values into a discrete value range using equation (5): Discrete value lower limit = $(e^{0.00198})^{10}$ = 1.02000 and the discrete value upper limit = $(e^{0.17022})^{10}$ = 5.48600. Thus, the 95% confidence interval for a $1 investment over a 10-year horizon with an annualized return of 10% and an annualized standard deviation of 15% would be between $1.02 and $5.49.

Although the mathematics may appear complicated, it is quite easy to program this into an Excel spreadsheet and immediately attain confidence intervals for any given mean, standard deviation, initial wealth value, and confidence level.[3] In addition, the model can easily be extended to incorporate investing throughout the time horizon. For example, if one wanted to calculate an expected value and a confidence interval for a 10-year annuity, one would simply sum up the expected values, lower limits, and upper limits over the first 10 years. This is based on the fact that the first payment has a 10 year horizon, the second payment has a nine year horizon, and so forth. This procedure will be accurate as long as returns from one year to the next are independent so that the expected values and limits are additive.

4. Theoretical versus bootstrapped values

Even with theoretically correct values, the question remains as to whether these forecast intervals calculated under the assumption of normality coincide with actual realized values. Although the bigger question is whether they *will* coincide with realized values in the future, that question is unanswerable. However, past results can be examined and possible future return paths can be simulated to see how well the theoretical projections would supposedly perform. This is how bootstrapping analysis works.

To determine the relationship between theoretical values and simulated values, simple equity/bond portfolios are formed ranging from 100% equity to 100% bonds. For equity returns, CRSP's value weighted index is used while bond returns are based on a constant 10-year T-bond maturity. For the all equity portfolio, calculations are based on data from January 1926 to December 2002. The mean and standard deviation for this time period is 11.66% and 21.23%, respectively. For the 50/50 stock/bond portfolio and the 100% bond portfolio, calculations are based on data from April 1954 to December 2002. The mean/standard deviation for these two portfolios are 9.71%/9.45% and 6.78%/6.80%, respectively.

Because there are only 80 years of equity data and 50 years of 10-year Treasury bond data, there are not enough independent samples that can be examined to validate the theoretical values. However, by employing bootstrapping that takes random samples with replacement from the actual returns that have been realized, thousands of possible return paths can be examined. In this way, no assumption about the distributional properties of asset returns need to be made and confidence intervals can be calculated.[4] Most of the studies discussed in Section 2 employ this procedure.[5]

Using 30,000 simulated samples, Table 2 shows the lower and upper level confidence intervals derived from this procedure for time periods from one to 40 years for an all equity portfolio, a 50/50 equity/bond portfolio, and a 100% bond portfolio.[6] As one can see when examining Table 2, the confidence intervals derived from bootstrapping closely coincide with the theoretical values that are calculated under the assumption of normality. As an example, the lower and upper annualized returns for an all equity portfolio associated with a 30 year time horizon are 2.54% and 17.35%, respectively, for the theoretical model. The bootstrapped values are almost identical, 2.49% and 17.44%, respectively. Although not reported, 90% and 99% confidence intervals are also tested to determine if extreme values may affect the validity of the normality assumption. The similarity between the theoretical and bootstrapped values remains the same.

In addition, the theoretical model can also be used to perform standard probability analysis by rearranging Eq. (4) and solving for Z. This value can then be used to determine the probability of ending with a value no less than some predetermined value. Furthermore, the theoretical model could be used to solve for the probability of reaching any value, which is partly the focus of Booth's (2004) study. As an example, by letting Z equal zero in Eq. (4), both the lower and upper limits would be equal. This value would be the minimum value that the portfolio would be worth 50% of the time. It is important to note as Booth correctly pointed out, that because of the lognormal nature of continuously compounded returns, this value is well short of the typical expected value.

As an example, for the 100% equity portfolio, after 30 years, the expected value using the discrete annualized average return of 11.61% is $27.02 for each dollar invested. However, there is only 30.3% chance of reaching this value. To actually attain a 50% probability of reaching a certain value, expectations would have to be reduced to $15.88 per $1 invested.

To further demonstrate the use of probability estimates, let's return to our hypothetical investor whose goal is to have $500,000 in 20 years for retirement. If the investor allocates all of his or her portfolio to equities that have an expected return of 11.66%, $500,000/

Table 2
95% confidence intervals from 1 to 40 years based on theoretical projections and 30,000 bootstrapped values for a 100% equity, 50/50 equity/bond, and a 100% bond portfolio. Lower and upper limits are shown on a per dollar basis with percentage difference between theoretical and bootstrapped values in parenthesis.

	Lower limit-annual return	Upper limit-annual return	Lower limit per $	Upper limit per $
100% Equity				
Theoretical				
1	−24.18%	58.71%	$0.76	$1.59
5	−7.01%	29.40%	$0.70	$3.63
10	−2.40%	23.29%	$0.78	$8.11
20	1.00%	19.14%	$1.22	$33.20
30	2.54%	17.35%	$2.12	$121.38
40	3.47%	16.29%	$3.92	$418.74
Boot strapping			(% difference)	(% difference)
1	−25.61%	57.87%	$.74 (−1.89%)	$1.58 (−.53%)
5	−7.33%	29.37%	$.68 (−1.70%)	$3.62 (−.12%)
10	−2.64%	23.33%	$.77 (−2.45%)	$8.14 (.37%)
20	0.82%	19.41%	$1.18 (−3.50%)	$34.73 (4.61%)
30	2.49%	17.44%	$2.09 (−1.39%)	$124.31 (2.41%)
40	3.34%	16.35%	$3.72 (−5.11%)	$426.73 (1.91%)
50/50 equity/bond				
Theoretical				
1	−7.64%	29.36%	$.92	$1.29
5	1.38%	17.86%	$1.07	$2.27
10	3.64%	15.29%	$1.43	$4.15
20	5.27%	13.50%	$2.79	$12.59
30	6.00%	12.72%	$5.74	$36.32
40	6.44%	12.26%	$12.12	$102.02
Boot strapping			(% difference)	(% difference)
1	−7.86%	29.01%	$0.92 (−.024%)	$1.29 (−.27%)
5	1.10%	17.56%	$1.06 (−1.36%)	$2.25 (−1.26%)
10	3.41%	15.12%	$1.40 (−2.21%)	$4.09 (−1.45%)
20	5.02%	13.36%	$2.66 (−4.55%	$12.27 (−2.57%)
30	5.75%	12.55%	$5.35 (−6.83%)	$34.72 (−4.41%)
40	6.21%	12.08%	$11.15 (−7.96%)	$95.88 (−6.02%)
100% Bond				
Theoretical				
1	−5.93%	20.72%	$.94	$1.21
5	.78%	12.68%	$1.04	$1.82
10	2.44%	10.85%	$1.27	$2.80
20	3.63%	9.58%	$2.04	$6.23
30	4.17%	9.02%	$3.40	$13.34
40	4.48%	8.69%	$5.78	$27.99
Boot strapping			(% difference)	(% difference)
1	−5.61%	21.21%	$.94 (0.33%)	$1.21 (.40%)
5	1.00%	12.75%	$1.05 (1.08%)	$1.82 (.34%)
10	2.56%	10.85%	$1.29 (1.15%)	$2.80 (.01%)
20	3.73%	9.62%	$2.08 (1.82%)	$6.28 (.79%)
30	4.26%	9.03%	$3.50 (2.83%)	$13.38 (.34%)
40	4.56%	8.66%	$5.95 (2.99%)	$27.76 (−.82%)

1.1166^{20} = \$55,083 is currently needed. Any shortfall would have to be met with additional contributions. However, if the investor is only willing to take a 2.5% chance that the portfolio value falls below the \$500,000 goal, $\$500,000/1.01^{20}$ = \$409,772 is currently needed. This figure is based on the lower annualized limit of 1% in Table 2 derived from the theoretical model.

Alternatively, if the investor is very concerned about the shortfall and cannot make up the difference, the results in Table 2 suggest more can be invested now in a less aggressive portfolio which will reduce the minimum amount needed to reach the \$500,0000 plateau with a 97.5% probability. As an example, if the investor decides on the 50/50 equity/bond mix, the expected return is reduced, but the lower annualized theoretical limit is higher for this portfolio over a 20-year horizon. In this example, only $\$500,000/1.0527^{20}$ = \$179,010 is needed to reduce the probability of failure to 2.5%. This may or may not be more attractive to the investor, but one can see the value of the knowledge gained by performing this type of analysis.

The more relevant point is that although the bootstrapped values will give virtually the same results, the theoretical model allows very quick adjustments to the expected return/standard deviation inputs as well as to the confidence level to evaluate a variety of situations. Virtually all one has to do is change the three inputs in an Excel spreadsheet. Furthermore, bootstrapping methods rely solely on historical data. Current market conditions may require adjustments to the historical expected returns and standard deviations to more accurately reflect future projections. To the extent that a more accurate forecast for expected returns and standard deviations could be made, theoretical models will actually be more accurate than bootstrapping techniques.

VAR analysis is also easily implemented using these models. In this case, the initial question would be how much could be lost given a particular probability. Although the application of VAR analysis for an extended forecast is complicated by the lognormality of continuous returns, it is still little more than a variation on a theme [see Kritzman (2003) if interested in the specifics]. Because standard confidence interval analysis can be accurately applied using theoretical models that assume normality, then VAR analysis can be similarly applied using the same models without needing to resort to bootstrapping to attain more accuracy. The results from the two procedures for the type of analysis above would be virtually the same.

As an example, standard VAR would simply be defined as the amount at risk below some set predefined value associated with a particular probability. Because Table 2 was created using a 95% confidence level, the value at risk below the initial \$1 investment associated with a 2.5% probability for any time frame is simply the value below \$1 in column four of Table 2. For the 100% equity portfolio, the VAR for a 10-year time horizon using the theoretical model is \$0.22 while VAR using bootstrapping is \$0.23.

It should also be pointed out that unless one is quite savvy with computer programming and knows all the "what if" scenarios that need to be answered, additional computer simulations would inevitably be needed as one studied the range of possibilities for a particular portfolio. On the other hand, theoretical methods allow the inclusion and analysis of extraordinary events that bootstrapping cannot incorporate if there is no historical precedence. Data could be artificially created to overcome this, but it is far easier to simply change the mean and standard deviation within the theoretical model to garner the results. Thus, for simple probabilistic projection analysis, theoretical models are actually more flexible in some ways and give virtually the same results as bootstrapping methods.

As a final note, the results of this study most assuredly extend to include portfolios consisting of asset mixes beyond simple equity/bond portfolios. This is because the lack of normality in equity returns, the possible lack of independence within equity returns, and certainly the lack of independence within bond returns and between equity and bond returns did not adversely affect the results. In fact, there is no reason to suspect that more complicated correlation structures associated with additional asset categories within a portfolio would cause sufficient non-normality for the portfolio as a whole to invalidate the results of this study. The opposite is more likely as additional asset categories would smooth portfolio returns and make the portfolio appear more, not less normal. Given this, for most investors' portfolios, once a relatively accurate estimate of a portfolio's mean and standard deviation is made, a theoretical probability projection can be confidently made.

5. Conclusions

Making the assumption that equity/bond portfolio returns are normally distributed appears to be valid when it comes to projecting confidence intervals over an extended period of time, or for determining probability estimates associated with certain terminal wealth values. Use of the standard approximation method should be applied with care as it is only relatively accurate up to a period of 10 years. Beyond 10 years, correcting for the lognormal nature of compounded returns is critical. After doing so, theoretically derived confidence interval projections coincide very closely with bootstrapped values which take the historical realization of returns as given. Given the results of this study, there is no reason to suspect that assuming normality to make long-term confidence interval projections would be any less accurate than bootstrapping methods for typical portfolios. Future uncertainty and the similarity of the values attained from the two procedures eliminate any possible accuracy advantage that bootstrapping methods have over theoretical models within the context of this study's analysis.

Thus, for simple "vanilla" forecasting, if we are armed with an expected mean and standard deviation for our portfolio, the theoretical model is just as effective as bootstrapping and far simpler to manipulate for a range of confidence interval projections, or for determining probabilities associated with various terminal values. As for more complicated portfolio structures, especially those that may exhibit strong non-normality, the greater robustness of bootstrapping is likely to take precedence.

However, since most investor's portfolios are composed of long positions in typical asset classes, applying the standard logarithmic model detailed in this study should yield accurate long-term probability forecasts. Bootstrapping procedures are not needed nor access to the large data sets to perform them. The only major inputs an investor or financial planner need be concerned with is the expected return and standard deviation of the portfolio. Assuming history is any guide, the expected return should range from 3% to around 10% as one progresses from 100% in a money market account to 100% in an all U.S. equity account.

Likewise, the standard deviation will range from 0% to approximately 20% for most investors. Historical estimates of expected returns and standard deviations could be used for a variety of tactical asset allocation configurations.

Alternatively, there are programs such as Morningstar's Principia that calculate an expected return and standard deviation for a portfolio based on the individual mutual funds that are held. Obviously these calculations can be erroneous and may or may not be relevant to future returns and volatility. However bootstrapping methods will not make the projections any more accurate and would require additional simulations for every conceivable portfolio mixture that is examined. The theoretical model does not suffer from this difficulty and can easily be used by any investor armed with an idea about the value of their portfolio's expected return and standard deviation.

Notes

1. Strictly speaking, all the studies discussed in this paper actually employ bootstrapping that is a special case of Monte Carlo analysis. Bootstrapping involves sampling an original data set with replacement to generate possible outcomes. Alternatively, Monte Carlo simulation is generally considered a procedure that generates possible outcomes by sampling from a theoretical distribution with predefined parameters.
2. To get an idea of how computationally intensive bootstrapping methods can be, a simple spreadsheet running 10,000 simulations at a time with just two asset categories, five portfolios, and six time frames took well over 30 minutes on an AMD Athlon 1.33 gigahertz processor.
3. The interested reader is welcome to download the spreadsheet I created which can be found at http://www.wku.edu/~bill.trainor/invest/riskandreturn.xls. Based on just four inputs: an expected return, standard deviation, confidence level, and beginning wealth level, the spreadsheet automatically calculates lower and upper wealth limits as well as lower and upper annualized returns for time horizons ranging from 1 to 40 years.
4. Note that this procedure assumes the returns are independent. To the extent that equity returns are mean reverting, this procedure could be biased. To combat this, block bootstrapping was also applied that take blocks of the data. Up to 2.5 year blocks were tested with only minor changes in the results, and are thus not reported. In general, the limits were reduced. As an example, for the 30 year 100% equity portfolio, the block bootstrapped annualized lower and upper limits were 1.38% and 17.37% compared to the basic bootstrapped values of 2.49% and 17.44%. These results held for the bond portfolios as well despite the fact that interest rates are well known to exhibit persistence.
5. Cooley et al. (2003) actually performed tests comparing bootstrapping methods to overlapping periods for estimating sustainable withdrawal rates in retirement. The reason the overlapping periods method does not work here is that for 40-year estimates, virtually the same time frame would have to be continuously sampled. Thus, this method is not discussed or applied for this type of analysis.
6. A 20/80 and an 80/20 equity/bond portfolio were also examined with similar results.

Acknowledgments

The author wishes to thank both Conrad Ciccotello (editor) and an anonymous reviewer for helpful comments and suggestions that substantially improved this paper in both content and form.

References

Affleck-Graves, J., & McDonald, B. (1989). Non-normalities and tests of asset-pricing theories. *Journal of Finance*, *44*, 889–908.

Bodie, Z., Kane, A., & Marcus, A. (2002). *Investments, fifth edition* (pp. 254–255). New York: McGraw Hill.

Booth, L.(2004). Formulating retirement targets and the impact of time horizon on asset allocation. *Financial Services Review*, *13*, 1–17.

Cooley, P. L., Hubbard, C. M., & Walz, D. (2003). Comparative analysis of retirement portfolio success rates: simulation versus overlapping periods. *Financial Services Review, 12,* 115–128.

Haim, S., & Yitzhaki, S. (2002). Estimating beta. *Review of Quantitative Finance and Accounting*, *18,* 95–118.

Hickman, K., Hunter, H., Byrd, J., & Terpening, W.(2001). Life cycle investing, holding periods and risk. *Journal of Portfolio Management*, *27,* 101–111.

Knez, P., & Ready, M. (1997). On the robustness of size and book-to-market in cross sectional regressions. *The Journal of Finance*, *52,* 1355–82.

Kritzman, M., & Chow, G. (2001). Value at risk for portfolios with short positions. *Journal of Portfolio Management*, *27,* 56–64.

Kritzman, M. (2003). *The portable financial analyst, second edition* (pp. 148–151). New York: Wiley.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, *36,* 394–419.

Milevsky, M. (2001). Spending your retirement in Monte Carlo. *Journal of Retirement Planning*, *4,* 21–29.

Mukherji, S. (2003). Optimal portfolios for different holding periods and target returns. *Financial Services Review*, *12,* 61–71.

Pye, G. B. (2000). Sustainable investment withdrawals. *Journal of Portfolio Management*, *26,* 73–83.

Samuleson, P. (1963). Risk and uncertainty: a fallacy of large numbers. *Scientia*, *98,* 108–113.

Vorkink, K. (2003). Return distributions and improved tests of asset pricing models. *The Review of Financial Studies*, *16,* 845–874.

Zhou, G. (1993). Asset pricing tests under alternative distributions. *Journal of Finance*, *48,* 1927–1942.