FAJ
Financial Analysts Journal
Volume 62 • Number 6
©2006, CFA Institute

The Misuse of Expected Returns

Eric Hughson, Michael Stutzer, and Chris Yung

Much textbook emphasis is placed on the mathematical notion of expected return and its historical estimate via an arithmetic average of past returns. But those wanting to forecast a typical future cumulative return should be more interested in estimating the median future cumulative return than in estimating the mathematical expected cumulative return. For that purpose, continuous compounding of the mathematical expected log gross return is more relevant than ordinary compounding of the mathematical expected gross return.

Popular finance textbooks and other methodological treatises emphasize the relevance of a portfolio's expected return and the use of time-averaged historical returns as an estimate of it. For example, Bodie, Kane, and Marcus (2004) state:

> . . . if our focus is on future performance then the arithmetic average is the statistic of interest because it is an unbiased estimate of an asset's future returns. (p. 865)

A more detailed procedure for using this average is found in a respected researcher's survey article:

> When returns are serially uncorrelated—that is, when one year's return is unrelated to the next year's return—the arithmetic average represents the best forecast of future return in any randomly selected year. For long holding periods, the best return forecast is the arithmetic average compounded up appropriately. (Campbell 2001, p. 3)

For an illustration of the quoted concepts, consider a hypothetical broad-based stock index with returns that are consistent with the ubiquitous random walk hypothesis. Annual gross (i.e., 1 plus net) returns for each of the past 30 years from the hypothetical index are given in the second column of **Table 1**. The arithmetic average gross return is 1.054 (i.e., the net return averages 5.4 a year). The last column provides the historical cumulative returns. The fourth column of Table 1 shows the cumulative return forecasts calculated by following the advice in the Campbell (2001) quotation—namely, compounding the arithmetic average to produce cumulative return forecasts at each future horizon between 1 and 30 years. For example, the forecasted cumulative return after 1 year is $1.054^1 = 1.054$, and after 30 years, it is $1.054^{30} = 4.799$; that is, an initial investment of $1.00 is forecasted to grow to about $4.80. But this return forecast is far higher than the 30-year historical cumulative return (3.005) shown at the bottom of the last column, which suggests that the arithmetic average–based forecast in the fourth column may be too high. We will now document that such overblown forecasts are very likely to happen in practice.

The overoptimism inherent when the arithmetic average return is used to forecast is illustrated in **Table 2** and **Figure 1**, which report the results from a bootstrap simulation of one million possible future cumulative returns derived from the annual gross returns given in Table 1.[1] Both Table 2 and Figure 1 clearly show that the mathematical expected cumulative return is always higher than the median cumulative return (i.e., the return that has equal chances of being exceeded or not) and that the gap between the two increases as the time horizon lengthens and the cumulative return distribution becomes more highly skewed to the right. For example, at the 10-year horizon, the mathematical expected cumulative return is 1.72, which is 18 percent higher than the median cumulative return (1.46). At the 30-year horizon, the mathematical expected cumulative return is 67 percent higher than the median cumulative return. As a result, the mathematical expected cumulative return is less likely to be realized (i.e., met or exceeded by the future cumulative return) in the future than the median return, and this likelihood is more pronounced for the long horizons used by retirement planners. For example, there is a 38 percent probability that the mathematical expected

Eric Hughson and Chris Yung are professors of finance in the Leeds School of Business at the University of Colorado, Boulder. Michael Stutzer is professor of finance and director of the Burridge Center for Securities Analysis and Valuation in the Leeds School of Business at the University of Colorado, Boulder.

Table 1. Forecasts Based on Historical Arithmetic Average Returns

Historical Period, T (years)	Gross Return	Log Gross Return	T-Year Cumulative Return Forecast	Historical Cumulative Return
1	1.014	0.014	1.054	1.014
2	0.876	–0.133	1.110	0.888
3	1.100	0.095	1.170	0.976
4	1.284	0.250	1.233	1.254
5	1.269	0.239	1.299	1.592
6	1.375	0.319	1.368	2.189
7	0.764	–0.269	1.442	1.673
8	1.024	0.024	1.519	1.713
9	1.250	0.223	1.601	2.141
10	0.901	–0.104	1.687	1.929
11	0.956	–0.045	1.777	1.845
12	0.823	–0.195	1.873	1.518
13	0.804	–0.218	1.973	1.221
14	0.916	–0.088	2.079	1.118
15	0.944	–0.057	2.191	1.056
16	0.772	–0.259	2.308	0.815
17	0.974	–0.026	2.432	0.794
18	0.998	–0.002	2.563	0.792
19	1.082	0.079	2.700	0.857
20	1.004	0.004	2.845	0.861
21	1.010	0.010	2.998	0.869
22	1.003	0.003	3.159	0.872
23	1.297	0.260	3.328	1.131
24	1.047	0.046	3.507	1.184
25	1.031	0.031	3.695	1.221
26	0.982	–0.018	3.893	1.199
27	1.426	0.355	4.102	1.709
28	1.208	0.189	4.323	2.064
29	1.515	0.415	4.555	3.126
30	0.961	–0.039	4.799	3.005
Average	1.054	0.037		

Table 2. Forecasting the Mathematical Expected and Median T-Year Cumulative Return

Horizon (T years)	Mathematical Expected Return	Compounded Arithmetic Average Return	Median Return	Compounded Average Log Return
5	1.31	1.30[a]	1.20	1.203[b]
10	1.72	1.69	1.46	1.45
20	3.01	2.86	2.18	2.10
30	5.15	4.84	3.09	3.03
40	9.01	8.20	4.72	4.39

[a] 1.054^5.

[b] $(e^{0.037})^5$.

Figure 1. Mathematical Expected Return vs. Median Cumulative Return: Bootstrap Simulation of Table 1 Returns

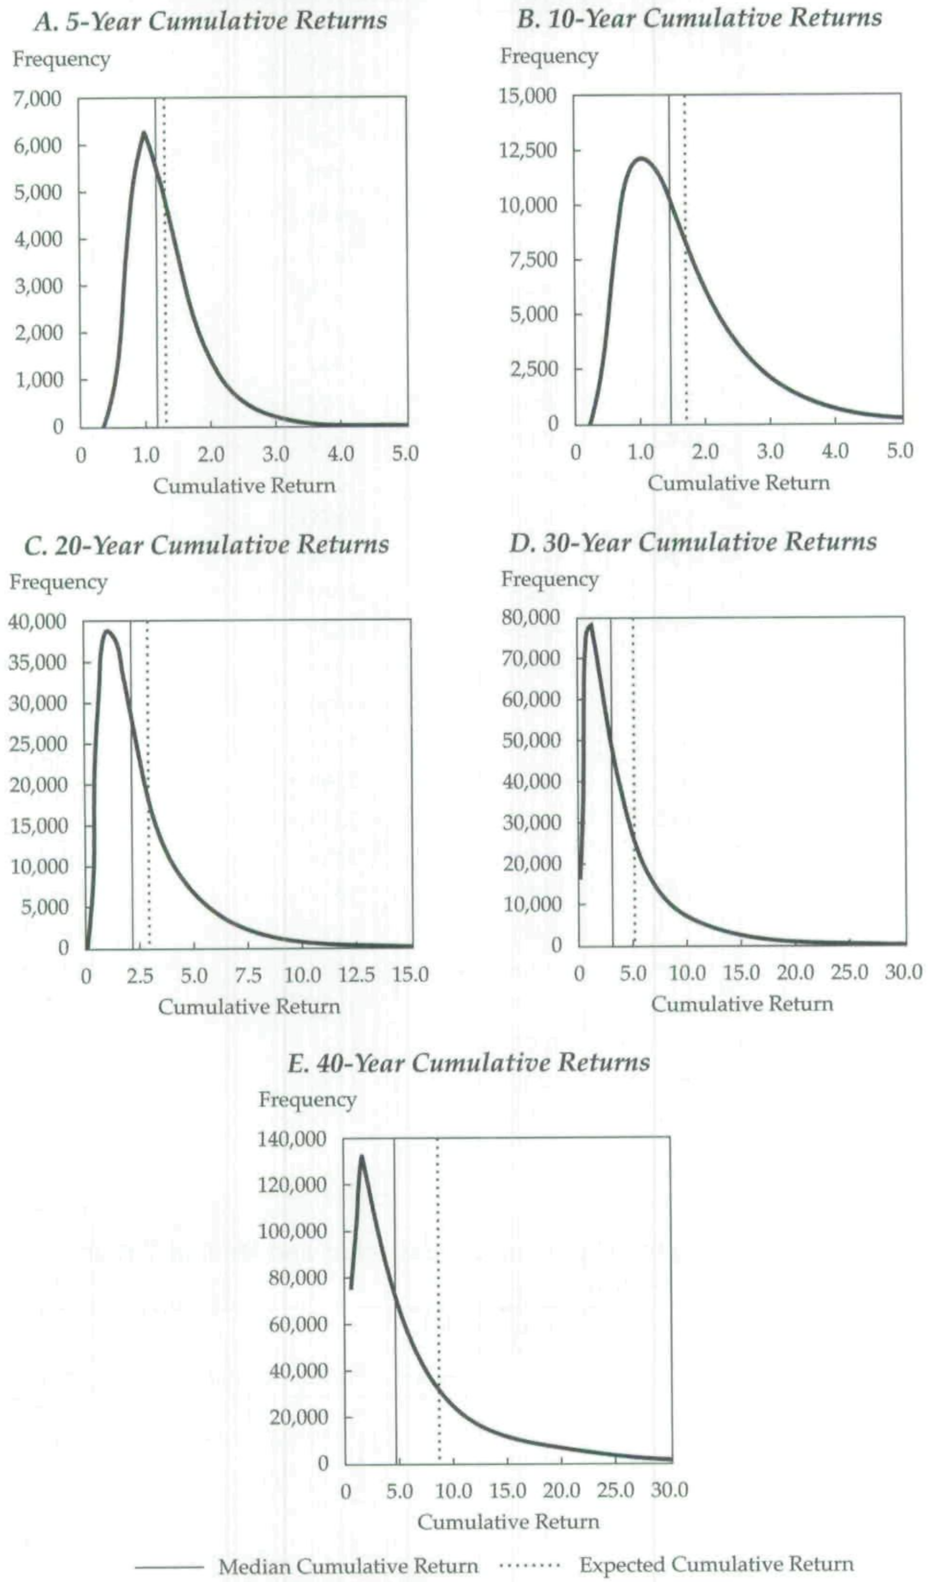

cumulative return will be exceeded at the 10-year horizon and only a 30 percent probability that it will be exceeded at the 30-year horizon.

The third column in Table 1 contains the logarithms of the historical gross returns. The average of these logarithms is lower than the arithmetic average of the gross returns themselves. Table 2 contrasts forecasts that compound the arithmetic average gross return with those that (continuously) compound the average *log* gross return (3.7 percent from Table 1). It is also common for analysts to call the number $e^{0.037} - 1 \approx 0.038$ percent the *geometric average* net return, in which case the last column of Table 2 is equivalently produced by ordinary compounding of the geometric average gross return (i.e., 1.038^T).[2] Table 2 shows that the compounded average of the log gross returns is far closer to the simulated median future cumulative return than is the compounded arithmetic average (1.054^T), which in turn, is far closer to the simulated mathematical expected future cumulative return. At the relatively long horizons that characterize retirement planning, the unwarranted optimism inherent in the arithmetic average–based forecasts will probably lead to excessively high investment in stocks.

To confirm that these problems also occur when actual monthly historical returns are used, we applied the same bootstrap simulation technique to the widely used 1926–2004 large-capitalization stock monthly returns produced by CRSP. The results, depicted in **Figure 2** and **Table 3**, confirm the previous problems. Moreover, a proof in Appendix A shows that these phenomena are generic, not simply the result of the specific data or accuracy of the simulations.

These findings are important because some investors do use the overly optimistic forecast procedure based on the historical arithmetic average. For example, in 2001, the chief actuary of the U.S. Social Security Administration described forecast procedures used in the organization's study of individual retirement account options that had been proposed but not yet enacted. The actuary noted that

> for individual account proposals, analysis of expected benefit levels and money's worth was also provided using a higher annual equity yield assumption of about 9.6 percent. This higher average yield reflected the arithmetic mean, rather than the geometric mean (which was 7 percent), of historical data for annual yields. (Campbell 2001, pp. 55–56)

In other words, the actuary made separate forecasts by compounding equity accounts at both the 9.6 percent historical arithmetic average return and the 7 percent geometric average return. We now examine possible reasons for compounding at the historical arithmetic average return rate.

Why Use the Arithmetic Average Return in Forecasting?

Two motives are put forth for using the arithmetic average return, but neither is convincing. The first motive is somewhat complex. Recall that the mathematical expectation of something is the probability-weighted average of its possible values. The quotations that began this article use this mathematical definition of expectation. When portfolio gross returns R_t are independently (I) and identically distributed (ID), the mathematical expected cumulative return (denoted by E) is the compounded value of the expected gross return per period—that is,

$$E(W_T) \overset{(\mathrm{I})}{=} \prod_{t=1}^{T} E(R_t) \overset{(\mathrm{ID})}{=} E(R_1)^T,$$

where W_T denotes the (random) cumulative return T periods in the future.

In addition, with the same IID assumption, the arithmetic average of historical gross returns is a commonly used estimate of the (unknown) constant mathematical expected gross return $E(R_1)$ per period, which becomes a more accurate estimate as the calendar history of gross returns lengthens. This argument is the typical motivation (as used in the opening quotation) for substituting the arithmetic average gross return (i.e., 1.054 percent) for the unobserved expected gross return per period, $E(R_1)$.[3] But as Figures 1 and 2 and Tables 2 and 3 show, the mathematical expected *cumulative* return, $E(W_T)$, for a stock index is less likely to be equaled or exceeded than the median cumulative return is. The situation becomes extreme in the long run because, as proven in Appendix A,

$$\text{prob}\left[W_T \geq E(W_T)\right] \overset{T\to\infty}{=} 0.$$

So, perhaps some long-term investors want to forecast the expected cumulative return because they use the word "expected" in its dictionary sense, rather than its mathematical sense. According to the Merriam-Webster Online Dictionary, "expect" means "to anticipate or look forward to the coming or occurrence of" or "to consider probable or certain." Long-term investors using this sense of the word will want to forecast the median cumulative return rather than the mathematical expected cumulative return because the unknown future cumulative return is more likely to equal or exceed the median. Tables 2 and 3 suggest that such investors should continuously compound the average *log* gross return (or, equivalently, simply compound the geometric average return) when making forecasts based solely on historical data.

Figure 2. Mathematical Expected vs. Median Cumulative Return: Bootstrap Simulation of (1926–2004) Large-Cap Stock Returns

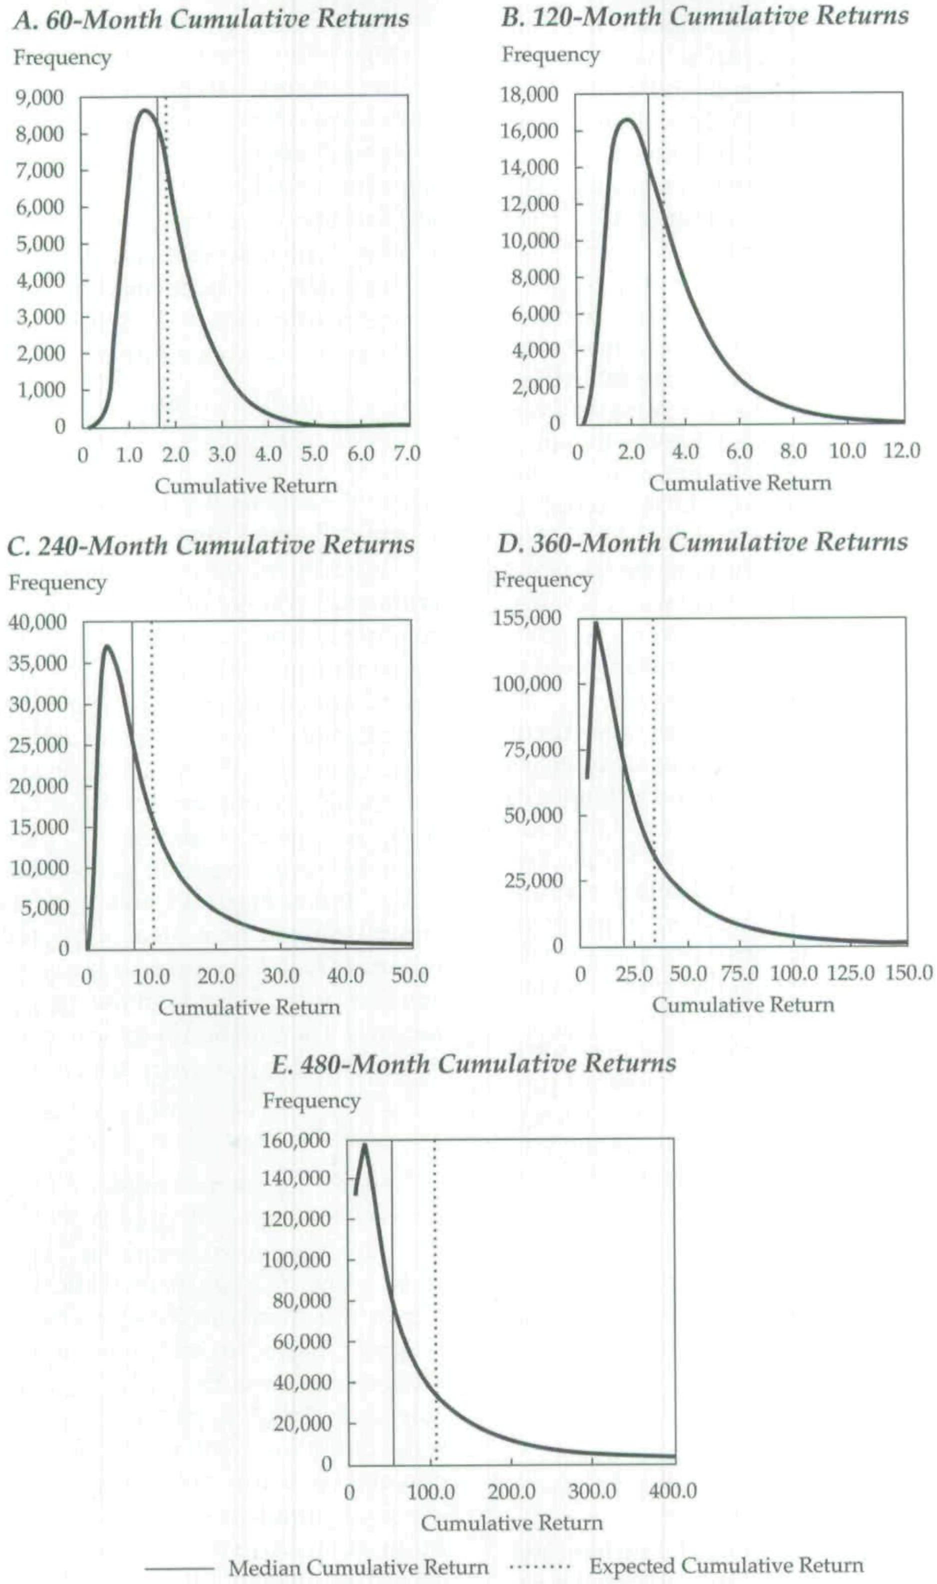

Table 3. Forecasting Based on Monthly (1926–2004) Large-Cap Returns

Horizon (T years)	Mathematical Expected Return	Compounded Arithmetic Average Return	Median Return	Compounded Average Log Return
5	1.79	1.82[a]	1.64	1.62[b]
10	3.22	3.30	2.68	2.64
20	10.35	10.89	7.16	6.96
30	33.26	35.95	19.08	18.38
40	106.69	118.65	50.83	48.51

[a] 1.010^{60}.
[b] $(e^{0.008})^{60}$.

A second possible motive for interest in forecasting the mathematical expected cumulative return arises from the statistical theory of best forecasts. This theory requires that the forecaster choose a *loss function* that quantifies the loss incurred by missing the forecast. Of course, the forecast error is random, so the best forecast is the one that minimizes a misforecasting "cost," defined to be the mathematical expected loss. When the loss is proportional to the *squared* forecast error, it is well known (see, for example, Zellner 1990) that the best forecast, *in this sense of the word "best,"* is the mathematical expected cumulative return, despite the fact that it will be higher than the median cumulative return. To understand why, consider the loss associated with underforecasting an unusually high cumulative return. The loss is extremely high because it is found by squaring the error between the high cumulative return and the (lower) forecast. For example, suppose the forecast error is +3. Then, the loss is 9, which is three times the size of the forecast error itself. Now, because cumulative returns are inherently positively skewed, the chance of underforecasting by an unusually large amount is greater than the chance of overforecasting. As a result, to minimize the chance of underforecasting by an unusually large amount, the forecast that minimizes the mathematical expected squared forecast error will be higher than the median. But this mathematical result is merely an alternative way of characterizing the behavior of someone who uses the expected cumulative return as a forecast. It is *not* a recommendation that investors use the squared forecast error to measure the loss from misforecasting. In fact, statisticians have also shown that if investors use the forecast error *itself* to measure the loss from misforecasting, the median is the statistically best forecast. After seeing the results in this article, we believe that most long-term investors would consider the median cumulative return to be a better forecast than the expected cumulative return. Hence, they act "as if" the forecast loss is the forecast error itself, rather than the squared error.

Limitations of Historical Returns

Unfortunately, using a historical average to estimate either the unknown expected gross return per year or expected log gross return per year requires, even under the ideal statistical circumstances embodied in the IID assumption, a very long calendar history of returns. Under the IID assumption, the estimate becomes progressively more accurate as the number of available past years' returns gets larger. But the convergence to the unknown true number is typically slow. Measuring returns more frequently (e.g., monthly or daily instead of annually) does absolutely no good.[4]

For an illustration of this point, note that the hypothetical annual stock returns in Table 1 (and used to produce Figure 1 and Table 2) were randomly sampled from a lognormal distribution with a volatility of 15 percent. Suppose we would like to be 95 percent confident that the historical average log gross annual return is within 400 bps of the (unknown) expected log gross return per year. Appendix A shows that we would need more than 54 years of past log gross returns to ensure this confidence level. And ±400 bps per year is probably too wide an uncertainty band for many financial planning purposes.

Moreover, even if we did have a long calendar history of returns, how likely is it that those returns would continue being generated by the same IID process? If the probability distribution of the measured returns changes over time, the compounding of historical averages will be very misleading, especially for short- or medium-term forecasts. Asness (2005) noted:

> When it comes to forecasting the future, especially when valuations (and thus historical returns) are at extremes, the answers we get from looking at simple historical averages are bunk. (p. 37)

This opinion may be excessively harsh, but we do feel that using historical returns to implement non-parametric forecast procedures (such as all the ones mentioned in this article) does not solve all problems inherent in the difficult task of forecasting cumulative returns.

Conclusion

Textbooks and other methodological sources may discuss the mathematical differences between historical arithmetic average and geometric average returns but may not adequately advise practitioners about the proper use of these concepts when forecasting future cumulative returns. Under ideal statistical assumptions, the historical arithmetic average gross return is an unbiased estimator of the mathematical expected gross return per period. As others have noted, compounding the mathematical expected gross return (but not the historical arithmetic average return) produces the mathematical expected *cumulative* return. But because cumulative returns are positively skewed, the mathematical expected cumulative return substantially overstates the future cumulative return that investors are likely to realize, and the problem grows worse as the horizon increases. Those seeking a more realistic forecast procedure can approximate the median cumulative return by continuously compounding the mathematical expected log gross return per period, using the historical average log gross return to estimate the expected log gross return. Without a hundred years or more of accurate returns to average, however, that procedure may still provide a highly inaccurate estimate—even if the return distribution does not change over time.

The authors wish to acknowledge Gitit Gur-Gershgorin for assistance with the simulations and Garland Durham for comments on the mathematics.

This article qualifies for 1 PD credit.

Appendix A. Derivation of Mathematical Claims

Using W_T to denote the cumulative return in time period T from a dollar invested initially and using R_t to denote the gross return at time t (i.e., 1 plus the net return), we express W_T as

$$W_T = \prod_{t=1}^{T} R_t. \tag{A1}$$

When the gross return process is IID, the mathematical expectation (denoted by E) is

$$E\left(W_T\right) \overset{\text{(I)}}{=} \prod_{t=1}^{T} E\left(R_t\right) \overset{\text{(ID)}}{=} E\left(R_1\right)^T, \tag{A2}$$

which shows that compounding expected portfolio gross returns produces the portfolio expected cumulative return, a fact underlying the Campbell (2001) quotation at the beginning of this article.

Representing Equation A1 and Equation A2 a bit differently will soon prove useful. To do so, we take the logarithm of both sides of Equation A2 and then reexponentiate to show that

$$E\left(W_T\right) = \left[e^{\log E(R)}\right]^T. \tag{A3}$$

Taking the logarithm of both sides of Equation A1, dividing and multiplying by T, and then reexponentiating shows that

$$W_T = \left(e^{\sum_{t=1}^{T} \log R_t / T}\right)^T. \tag{A4}$$

We see from Equation A4 that it is the time-averaged *log* gross returns that determine the evolution of a portfolio's cumulative return, regardless of whether or not the returns are IID. Because the (nonlog) gross returns are never negative for stock and/or bond investments, and raising something nonnegative to a fixed power greater than 1 is a monotone nondecreasing function, Equation A3 and Equation A4 imply that the probability of doing at least as well as the expected cumulative return is

$$\text{prob}\left[W_T \geq E\left(W_T\right)\right] = \text{prob}\left[\frac{\sum_{t=1}^{T} \log R_t}{T} \geq \log E\left(R\right)\right]. \tag{A5}$$

But by the law of large numbers for IID processes,

$$\frac{1}{T}\sum_{t=1}^{T} \log R_t \overset{T\to\infty}{=} E\left(\log R\right). \tag{A6}$$

So, from Equation A5 and Equation A6, we see that the long-run behavior of the probability (Equation A5) is governed by the relationship between $E(\log R)$ and $\log E(R)$. Because Jensen's inequality implies that

$$E\left(\log R\right) < \log E\left(R\right), \tag{A7}$$

Equations A5–A7 imply the distribution-free result in the text; that is,

$$\text{prob}\left[W_T \geq E\left(W_T\right)\right] \overset{T\to\infty}{=} 0. \tag{A8}$$

From Equation A8, we can clearly see that we should not expect to earn a long-run cumulative return that is greater than or equal to the expected cumulative return! When a variable is not symmetrically distributed, its mathematical expectation is not generally a good indicator of the variable's central tendency. Figures 1 and 2 show that the cumulative return distribution is sharply skewed to the right, so the expected cumulative return, $E(W_T)$, is higher than what will likely occur.

The horizon-dependent probabilities (Equation A5) are easily calculated when the returns are lognormally distributed, as the hypothetical returns used in Table 1 are. Substituting our notation for that used by Hull (1993, p. 211), $\log W_T$ is normally distributed with mean equal to $(\mu - \sigma^2/2)T$ and variance equal to $\sigma^2 T$, where μ and σ are, respectively, the annualized mean and volatility parameters. Hull also showed that

$$E(W_T) = \left(e^{\mu T}\right) \equiv \left(e^{\mu}\right)^T \tag{A9}$$

(that is, the compound value of the expected gross return). Similarly,

$$E(\log W_T) = (\mu - \sigma^2/2)T. \tag{A10}$$

Because the logarithm is a monotone increasing transformation,

$$\begin{aligned} \text{prob}\left[W_T \geq E(W_T)\right] &\equiv \text{prob}\left(\log W_T \geq \log e^{\mu T}\right) \\ &= \text{prob}\left[Z \geq \frac{\mu T - \left(\mu - \sigma^2/2\right)T}{\sigma\sqrt{T}}\right] \\ &= \text{prob}\left(Z \geq \frac{\sigma}{2}\sqrt{T}\right), \end{aligned} \tag{A11}$$

where Z denotes the standard normal density function. We see that $\text{prob}[W_T \geq E(W_T)]$ approaches zero as the horizon, T, approaches infinity, as proven for the arbitrary distributions.

Fortunately, $\text{prob}[W_T \geq Median(W_T)] = 1/2$, instead of approaching 0 for large T. When the Hull (1993, p. 211) lognormal example is used, the median cumulative return is $e^{(\mu - \sigma^2/2)T}$; that is, it is produced by compounding the expected log gross return per year. The percentage difference of the expected and median cumulative returns is

$$\frac{\left(e^{\mu}\right)^T - \left[e^{\left(\mu - \sigma^2/2\right)}\right]^T}{\left[e^{\left(\mu - \sigma^2/2\right)}\right]^T} = e^{(\sigma^2/2)T} - 1, \tag{A12}$$

which is an increasing function of σ and T.

A particularly stark example of the difference between the expected and median cumulative returns can be seen by considering a volatile investment (e.g., $\mu = 8$ percent and $\sigma = 40$ percent). Then, the median cumulative return is

$$e^{\left(\mu - \sigma^2/2\right)T} = 1 \tag{A13}$$

for all horizons T (i.e., there would be no tendency for the investment value to drift either up or down, despite the seemingly high expected cumulative return $e^{0.08T}$). An investment with a smaller μ would result in negative drift (i.e., a tendency to lose money). However, a more diversified portfolio with a smaller return ($\mu < 8$ percent) would tend to make money if its volatility were low enough to make $\mu - \sigma^2/2 > 0$.

But what about when $\log R$ is not normal? We will now see why compounding the average log return produces a reasonable estimate of the median cumulative return whether the IID distribution of $\log R$ is normal or not. First, we rewrite Equation A4 by canceling T to obtain

$$W_T = e^{\sum_{t=1}^{T} \log R_t}. \tag{A14}$$

Because the exponential function is a monotone (increasing) function of $\sum_{t=1}^{T} \log R_t$,

$$Median(W_T) = e^{Median\left(\sum_{t=1}^{T} \log R_t\right)}. \tag{A15}$$

When T is suitably large, Ethier (2004) used the following approximation:

$$\begin{aligned} Median\left(\sum_{t=1}^{T} \log R_t\right) &\overset{\text{IID}}{\approx} E(\log R)T \\ &- \frac{E\left\{\left[\log R - E(\log R)\right]^3\right\}}{6 \operatorname{var}(\log R)}. \end{aligned} \tag{A16}$$

In practice, the second term in Equation A16 is quite small compared with the first term. So, substituting Equation A16 into Equation A15 yields

$$Median(W_T) \approx \left[e^{E(\log R)}\right]^T. \tag{A17}$$

It is easy to show the equivalence between compounding the historical average log gross return and compounding the historical geometric average net return. Denote the actual historical returns by $R_1^h, \ldots, R_N^h$ and the historical cumulative return by W^h. Then, the historical geometric average, $R(g)$, is defined by

$$W^h = \prod_{t=1}^{N} R_t^h \equiv \left[1 + R(g)\right]^N. \tag{A18}$$

But W^h can also be computed by

$$W^h = e^{\frac{\sum_{t=1}^{N} \log R_t^h}{N} N}. \tag{A19}$$

Equation A18 and Equation A19 show that

$$1 + R(g) = e^{\frac{\sum_{t=1}^{N} \log R_t^h}{N}}, \tag{A20}$$

so raising either side of Equation A20 to the power T produces the same T-period cumulative return forecast—a "plug-in" estimator of Equation A17. If $\tilde{R}_t^h$ is used to denote a generic random historical gross return, a desirable property of this plug-in estimator is that

$$Median\left(e^{\frac{\sum_{t=1}^{N}\log\tilde{R}_t^h}{N}}\right)^T = \left(e^{Median\,\frac{\sum_{t=1}^{N}\log\tilde{R}_t^h}{N}}\right)^T, \quad \text{(A21)}$$

which we dub "median unbiasedness." A more complex procedure might provide a better estimator of the median cumulative return, but simplicity of implementation and motivation are practitioners' desiderata that would be implicitly ignored by those (if any) who advocated a more complex procedure.

Unfortunately, historical arithmetic or geometric averages are inherently imprecise estimators—a fact that is easily illustrated under lognormality. The log of the one-year cumulative return distribution has an expected value of $\mu - \sigma^2/2$. Suppose we measure log returns $1/\Delta t$ times per year (e.g., $\Delta t = 1/12$ when returns are measured monthly). Then, the log gross return per measurement period has an expected value of $(\mu - \sigma^2/2)\Delta t$, so a historical average of $N \equiv T/\Delta t$ log gross returns (i.e., T years of history) will also be normally distributed with an expected value equal to $(\mu - \sigma^2/2)\Delta t$. Hence, an unbiased estimator of $\mu - \sigma^2/2$ is the historical average log gross return divided by Δt. Because the log gross return per measurement period has a variance of $\sigma^2\Delta t$, the variance of the unbiased estimator is $\sigma^2(\Delta t)^2/T$ divided by $(\Delta t)^2$, which equals σ^2/T and is hence independent of the return measurement interval, Δt. A 95 percent confidence interval for the historical average will then have a width of $\pm 1.96\sigma/\sqrt{T}$. For that width to be ± 0.04, T would have to be $\sigma^2(1.96/0.04)^2$ years. Substituting $\sigma = 0.15$ yields 54 years, no matter how frequently returns are measured, as claimed near the end of the text.

Notes

1. T random draws from the annual gross returns in Table 1 were multiplied to produce a possible T-year cumulative return. The procedure was repeated one million times to produce each of the smoothed histograms in Figure 1, whose means and medians are reported in Table 2.
2. See Appendix A for a derivation of this equivalence and the other mathematical claims made later in the text.
3. Even this typical motivation is flawed. An interesting article by Jacquier, Kane, and Marcus (2003) highlighted problems resulting from compounding the historical arithmetic average to produce a data-based estimate of the unknown mathematical expected cumulative return. They added the assumption that returns are lognormally distributed and proposed better estimators of the unknown mathematical expected cumulative return. Our goal is different, however, for we are highlighting flaws in the arguments used to justify the relevance of estimating the mathematical expected cumulative return in the first place. We argue that estimating the *median* cumulative return is a much more relevant objective, regardless of whether the returns are lognormal.
4. See Luenberger (1998) for a simple exposition of this problem and Appendix A for a specific calculation.

References

Asness, Clifford S. 2005. "Rubble Logic: What Did We Learn from the Great Stock Market Bubble?" *Financial Analysts Journal*, vol. 61, no. 6 (November/December):36–54.

Bodie, Zvi, Alex Kane, and Alan Marcus. 2004. *Essentials of Investments*. 6th ed. New York: McGraw-Hill.

Campbell, John Y. 2001. "Forecasting U.S. Equity Returns in the 21st Century." In *Estimating the Real Rate of Return on Stocks over the Long Term*. Edited by John Y. Campbell, Peter A. Diamond, and John B. Shoven. Social Security Advisory Board (www.ssab.gov/publications.htm).

Ethier, S.N. 2004. "The Kelly System Maximizes Median Fortune." *Journal of Applied Probability*, vol. 41, no. 4 (December):1230–1236.

Hull, John. 1993. *Options, Futures, and Other Derivative Securities*. Upper Saddle River, NJ: Prentice Hall.

Jacquier, Eric, Alex Kane, and Alan Marcus. 2003. "Geometric Mean or Arithmetic Mean: A Reconsideration." *Financial Analysts Journal*, vol. 59, no. 6 (November/December):46–53.

Luenberger, David G. 1998. *Investment Science*. New York: Oxford University Press.

Zellner, Arnold. 1990. "Bayesian Inference." In *The New Palgrave: Time Series and Statistics*. Edited by John Eatwell, Murray Milgate, and Peter Newman. New York: Norton.