Geometric or Arithmetic Mean: A Reconsideration

Eric Jacquier
Alex Kane
Alan J. Marcus

December 18, 2002
This revision: May 20, 2003

Jacquier is at Boston College, Wallace E. Carroll School of Management, HEC Montreal. (514) 340-6194; eric.jacquier@hec.ca and CIRANO, www.cirano.qc.ca. Kane is at University of California, San Diego, Graduate School of International Relations and Pacific Studies. (858) 534-5969; akane@ucsd.edu. Marcus is at Boston College, Wallace E. Carroll School of Management. (617) 552-2767; alan.marcus@bc.edu.

1

Geometric or Arithmetic Mean: A Reconsideration

Mini-abstract

Average past performance paints too rosy a view of future portfolio growth.

Geometric or Arithmetic Mean: A Reconsideration

Abstract

An unbiased forecast of the terminal value of a portfolio requires compounding its initial value at its *true* arithmetic mean return for the length of the investment period. However, compounding at the arithmetic average historical return results in an upward biased forecast. This bias does *not* necessarily disappear even if the sample average return is an unbiased estimator of the true mean, computed from long data series, and returns are generated by a stable distribution. On the other hand, forecasts obtained by compounding at the geometric average generally will be *downward* biased. For typical investment horizons, the proper compounding rate is in-between these two values. These biases are empirically significant. For investment horizons of 40 years, the difference in forecasts of cumulative performance easily can exceed a factor of two. The percentage differences in forecasts grows with the investment horizon, as well as with the imprecision in the estimate of the mean return.

Geometric or Arithmetic Mean: A Reconsideration

Digest

It is well known that an unbiased forecast of the terminal value of a portfolio requires the initial value to be compounded at the arithmetic mean rate of return for the length of the investment period. However, an upward bias in forecasted values results if one estimates the mean return with the sample average, and uses that average to compound forward. This bias results from the facts that cumulative performance is a non-linear function of average return, and that the sample average is necessarily a noisy estimate of population mean. Surprisingly, the bias does *not* necessarily disappear asymptotically, even if the sample average is computed from long data series and returns come from a stable distribution with no serial correlation. Instead, the bias depends on the *ratio* of the length of the historic estimation period to that of the forecast period.

On the other hand, forecasts obtained by compounding at the geometric average generally will be *downward* biased. For typical investment horizons, the proper compounding rate is in-between these two values. Specifically, unbiased estimates of future portfolio value require that the current value be compounded forward at a weighted average of the two rates. The proper weight on the geometric average equals the ratio of the investment horizon to the sample estimation period. Therefore, for short investment horizons, the arithmetic average will be close to the "unbiased compounding rate." However, as the horizon approaches the length of the estimation period, the weight on the geometric average approaches 1. For even longer horizons, both the geometric and arithmetic average forecasts will be upward biased.

4

The implications of these results are generally sobering. There is already an emerging consensus that the 1926-2002 historical average return on broad market indexes such as the S&P 500 are probably higher than likely future performance. Our results imply that the best forecasts of compound growth rates for future investments are even lower than the estimates emerging from this new strain of research.

The impact of these considerations on forecasts of future portfolio value can be dramatic. Compounding at the arithmetic average return calculated from sample periods of either the most recent 77 or 52 years results in forecasts of future value for a sample of countries that are roughly double the corresponding unbiased forecasts based on the same data periods. Indeed, for reasonable risk and return parameters, at investment horizons of 40 years, the difference in forecasts of total return generally will exceed a factor of two.

The percentage differences between unbiased forecasts versus forecasts obtained by compounding arithmetic or geometric average returns increase with the ratio of the investment horizon to the sample estimation period, as well as with the imprecision in the estimate of the mean return. For this reason, emerging markets present the greatest problem. These markets present historical estimation periods that are particularly short, and return histories that are particularly noisy. For these markets, therefore, the biases analyzed in this paper can be especially acute. Even for developed economies, which offer longer histories, there can be significant bias if one disregards data from very early periods because it is considered less predictive of future performance.

Geometric or Arithmetic Mean: A Reconsideration

Increased concern for long-term retirement planning, the associated growth of the defined-contribution investments market, and proposals for Social Security reform have all focused considerable attention on forecasts of long-term portfolio returns. Moreover, recent academic studies[1] suggest that conventional estimates of long-term performance, such as those guided by historical averages from the database in the popular Ibbotson Associates publication *Stocks, Bonds, Bills and Inflation Annual Yearbook* may paint far too rosy pictures of likely future performance.

This paper returns to an old controversy in the forecasting of long-term portfolio performance: given a historical data series from which one estimates the mean and variance of portfolio returns, should one forecast future performance using arithmetic or geometric averages? Finance texts[2] generally (and correctly) note that if the arithmetic mean of the portfolio's stochastic rate of return is known, an unbiased estimate of cumulative return is obtained by compounding at that rate. Despite this advice, many in the practitioner community prefer geometric averages, which are necessarily lower than arithmetic averages.

---

[1] Fama and French (2002) and Jagannathan, McGratten and Scherbina (2000) show that estimates of expected return derived from a dividend discount model are substantially lower than historical average returns. This suggests that U.S. experience has turned out better than market participants expected, rendering historical averages greater than equilibrium risk premia. In addition, more inclusive data bases result in lower historical risk premia. See for example, Dimson, Marsh and Staunton (2002), which shows that extending U.S. data to earlier historical periods reduces historical average returns.

[2] Including those of two of the authors; see for example, Bodie, Kane and Marcus (2002, pp. 810-811), but also Brealey and Myers (2003, pp. 156-157) or Ross, Westerfield, and Jaffe (2002, pp. 232-233).

We show in this paper that the practitioners are on to something. When the true mean return of the portfolio is not known, but must be estimated, we demonstrate that compounding at the sample arithmetic average return results in an upwardly biased estimate of final portfolio value *even if the sample estimate of mean return is itself unbiased*. Neither does the geometric average generally provide an unbiased forecast of cumulative return. But compounding at a simple weighted average of the arithmetic and geometric averages does provide an unbiased forecast of cumulative return.

The proper weights applied to the two averages depend on the ratio of the investment horizon to the length of the sample period used to estimate the mean return. When the forecast horizon is very short relative to the sample period, the weight on the arithmetic average return is close to 1.0; in this case, the "textbook wisdom" is correct. However, as the horizon extends, the weight progressively shifts in favor of the geometric average. At a horizon equal to the historical sample period, the weight on the geometric average reaches 1.0, and "practitioner wisdom" is correct. At even longer horizons, the weight on the geometric average exceeds 1.0, while that on the arithmetic average is negative, implying unbiased compounding rates below *both* the arithmetic and geometric averages.

The implications of these results are quite sobering. There is already an emerging consensus (see footnote 1) that the 1926-2002 historical average returns on broad market indexes such as the S&P 500 are probably higher than likely future performance. We conclude that the best forecasts of future compound growth rates are even lower than the estimates emerging from this new strain of research.

7

In the next section we show why unbiased estimators of mean return result in overestimates of compound portfolio growth rates. We illustrate there that the bias is a function both of the imprecision of the estimate and of the horizon for which portfolio performance is forecasted. We present, easy to use, exact formulas for this bias in the standard case that returns are log-normally distributed. The formula indicates a simple correction for bias, which involves the weighted average of the arithmetic and geometric average. In Section 2, we examine some implications of our results and demonstrate that this bias can be substantial, certainly large enough for significant implications for long-term planning. We show that the bias tends to be more severe when there is a shorter history of returns or when returns are more variable, making these issues particularly important for emerging markets. Empirical estimates in this section demonstrate that potential bias in these markets is enormous. Section 3 concludes

1. Forecasting Cumulative Returns with Noisy Estimates

Suppose the rate of return on a stock portfolio is log-normally distributed. If the stock price today, at time t, is denoted S_t, then $ln(S_{t+1}/S_t)$ has a normal distribution with mean μ and variance σ^2. Over an investment horizon of H periods, if returns are independent from one period to another, the cumulative return on the portfolio will also be log-normally distributed: $ln(S_{t+H}/S_t)$ has a normal distribution with mean μH and variance $\sigma^2 H$.

For any historical sample of stock returns, the geometric average rate of return is defined as the compound growth rate of portfolio value over the investment period.[3] Suppose for example that we have observed stock prices over a sample period starting T periods ago (i.e., starting at time $t - T$) and ending today, at time t. If the initial value of the portfolio was S_{t-T}, then the geometric average rate of return, g, is defined by

$$S_{t-T}\, e^{gT} \;=\; S_t$$

or equivalently,

$$g \;=\; \frac{1}{T}\, ln(S_t/S_{t-T}) \tag{1}$$

Because in each period, the expected value of $ln(S_{t+1}/S_t)$ equals μ, the geometric average return is an unbiased estimator (in fact, the maximum likelihood estimator) of μ.

However, it is a well-known feature of the lognormal distribution that if $ln(S_{t+1}/S_t)$ has mean μ, then the expected value of S_{t+1} equals $S_t\, e^{\mu + \frac{1}{2}\sigma^2}$. Thus, the expected rate of growth in portfolio value expressed at a continuously compounded rate is $\mu + \frac{1}{2}\,\sigma^2$. This quantity is the arithmetic mean rate of return, which exceeds the geometric mean by $\frac{1}{2}\sigma^2$. After an investment horizon of H periods, the unbiased forecast of future portfolio value is therefore

$$E(S_{t+H}) \;=\; S_t\, e^{(\mu + \frac{1}{2}\sigma^2)H} \tag{2}$$

Equation (2) is the basis of the "textbook rule" that to forecast future value, one should compound forward at the mean arithmetic return.

---

[3] For expositional simplicity, we assume the stock portfolio pays no dividends. If there were dividends, we would simply add the dividend yield to obtain the total rate of return.

The difference in these approaches can be empirically significant. We estimate the arithmetic mean by computing the growth in portfolio value each period (i.e., S_{t+1}/S_t) and then calculating the sample period average. This average is the estimate of $S_t\,e^{\mu + \frac{1}{2}\sigma^2}$. We estimate the geometric mean from equation (1). Using the *SBBI* database from 1926 to 2001, we find the geometric average annual return for the S&P 500 (expressed as a continuously compound rate) is 10.51%, while the arithmetic average is 12.49%. The standard deviation of the index over this period was 20.3%, or .203, so the difference in the two measures returns is just about half the variance ($\frac{1}{2} \times .203^2 = .0206$ or 2.06%), consistent with the fact that the annual return of the index is approximately log-normal. For more volatile investments such as small stocks, the difference in arithmetic and geometric average is even larger.

However, an often-overlooked presumption of the textbook formula is that the forecaster knows the true values of the parameters μ and σ. In practice, of course, these are estimated, and even using the best estimation techniques, the estimators are subject to sampling error.

One might think that simply substituting unbiased estimates of μ and σ into equation (2) would provide unbiased estimates of future portfolio value. In fact, this is common practice. For example, Ibbotson Associates simulates future portfolio values in *SBBI* by compounding forward using the historical arithmetic average as in equation (2). Unfortunately, even if the estimate of μ is unbiased and σ is known[4], the forecast of

---

[4] We will focus on the estimation error in μ, ignoring possible sampling error in σ. The justification for this simplification is that if the return distribution is stable, one can estimate σ arbitrarily accurately by sampling returns more frequently. In contrast, the precision of the estimate of μ depends solely on the

future portfolio value will be biased, possibly quite severely. The reason for this result is that $e^{(\mu + \frac{1}{2}\sigma^2)H}$ is a non-linear function of μ. Symmetric errors in the estimate of μ therefore have asymmetric effects on the forecast of $S_t\, e^{(\mu + \frac{1}{2}\sigma^2)H}$. Positive estimation error has a greater impact than an equal-magnitude negative error. Therefore, even if the estimate of μ is unbiased, with estimation error centered around zero, the estimation error in S_{t+H} will be upward biased.

Figure 1 illustrates this property. Suppose the true mean value of μ is 10%, the standard deviation of annual returns is 20%, and we estimate μ from equation (1) using returns over a 30-year period. The standard error of the estimate $\hat{\mu}$ is then $20/\sqrt{30} =$ 3.65%. Panel A (top left) shows that the probability density of $\hat{\mu}$ is symmetrically distributed around 10% with a standard deviation of 3.65%. However, the probability densities for forecasted final wealth (all based on an initial investment of $1) are right-skewed. Observe the dashed vertical lines in Panels B – D, all corresponding to forecasts of final portfolio value obtained using estimates equal to $\mu = 10\%$ plus or minus 3.65%. For short investment horizons such as 2 years (Panel B, bottom left) the effect of skewness is minimal. But for a 10-year horizon (Panel C, top right), a one-standard-error positive error in the estimate of μ increases the forecast of final value by $1.5 from $3.3 to $4.8, whereas the symmetric one-standard-error negative error in the estimate of μ reduces the forecast of final value by only $1.0, from $3.3 to $2.3. The asymmetry at a 20-year horizon is even more dramatic (Panel D, bottom right). In all cases, the uncertainty in final wealth is considerable.

---

length of the sampling period and cannot be enhanced by sampling more frequently. Merton (1980) demonstrates this result rigorously.

11

[Figure 1 here]

If the underlying stock price process is log-normal, it is relatively easy to derive the exact bias in the forecast.[5] If we estimate μ using equation (1), then our estimate, $\hat{\mu}$, equals the geometric average return over the sample period of length T. The standard error of $\hat{\mu}$ is $\sigma/\sqrt{T}$. The (noisy) forecast extends for H periods, resulting in a standard deviation of the forecast equal to $\sigma H/\sqrt{T}$, and variance $\sigma^2 H^2/T$. Thus, estimation error in μ increases the range of possible values we may infer for final portfolio value: in addition to the "irreducible noise" due to economic uncertainty (measured by σ), we add *additional* noise by using an estimate of μ to forecast. We know from equation (2) that adding variance to a log-normal return increases the forecast of cumulative portfolio growth by one-half the variance of cumulative return. Hence, the upward bias resulting from the extra volatility associated with sampling error is $e^{\frac{1}{2}\sigma^2 H^2/T}$. Note that the bias increases in both investment horizon, H, and volatility, σ (which will make our statistical estimates less precise). Conversely, bias declines with T, because longer sample periods increase the precision of our estimates.

Table 1 computes this bias as a function of investment horizon, volatility, and sample estimation period. The table demonstrates that using reasonable parameters, the bias can be dramatic, especially when volatility is high, or the sample period is short. In these cases, 30 or 40-year forecasts can be biased by factors of 2 or more.

---

[5] A more formal derivation of this formula for bias may be found in Eric Jacquier, Alex Kane, Alan J. Marcus, "Optimal Forecasts of Long-Term Returns and Asset Allocation: Geometric, Arithmetic, or Other Means?" working paper, 2002. Blume (1974) first discusses the bias in forecasted portfolio value. However, because he assumes that returns are normally rather than log-normally distributed, he does not obtain exact formulas for expected values or bias. A log-normal specification is also preferable as it rules out returns below −100%.

[Table 1 here]

We conclude that while the expected future value at horizon $t + H$ of a portfolio currently worth \$1 can be described by equation (2), we may not simply substitute an estimate of μ such as the historical geometric average into this formula. As noted, substituting $\hat{\mu}$ for μ adds extra variability to the distribution of portfolio values, and results in the following bias:

$$E[e^{(\hat{\mu} + \frac{1}{2}\sigma^2)H}] = E(S_{t+H})\, e^{\frac{1}{2}\sigma^2 H^2/T} \tag{3}$$

Equation (3) does, however, indicate how one can adjust the estimate of the compound growth rate of the portfolio to render the forecast of portfolio value unbiased. Suppose we start with the sample estimate of the continuously compounded arithmetic average rate of return, i.e., $\hat{\mu} + \frac{1}{2}\sigma^2$, but then reduce this estimate by the amount $\frac{1}{2}\sigma^2 H/T$. Call this modified estimator $\mu^* + \frac{1}{2}\sigma^2$. This reduction is just sufficient to undo the bias associated with the use of $\hat{\mu}$:

$$
\begin{aligned}
e^{(\mu^* + \frac{1}{2}\sigma^2)H} \;&=\; e^{(\hat{\mu} + \frac{1}{2}\sigma^2 - \frac{1}{2}\sigma^2 H/T)H} \;=\; e^{(\hat{\mu} + \frac{1}{2}\sigma^2)H}\, e^{-\frac{1}{2}\sigma^2 H^2/T} \\[6pt]
&=\; [E(S_{t+H})\, e^{\frac{1}{2}\sigma^2 H^2/T}] \times e^{-\frac{1}{2}\sigma^2 H^2/T} \\[6pt]
&=\; E(S_{t+H}) \tag{5}
\end{aligned}
$$

Thus $\mu^* + \frac{1}{2}\sigma^2$ is the compound growth rate that provides unbiased estimates of future portfolio value.

13

Now notice that this growth rate is a *weighted average* of the geometric and arithmetic averages with weights that depend on the ratio of the investment horizon to the sample estimation period[6]:

$$\mu^* + \tfrac{1}{2}\,\sigma^2 \equiv \hat{\mu} + \tfrac{1}{2}\,\sigma^2 - \tfrac{1}{2}\,\sigma^2 H/T = (\hat{\mu} + \tfrac{1}{2}\,\sigma^2)(1 - H/T) + \hat{\mu}\,H/T \qquad (6)$$

The growth rate that gives an unbiased forecast of final portfolio value will be very close to the arithmetic average for very short investment horizons (i.e., for which H/T is close to zero). But as the horizon extends, the weight on the geometric average increases. Therefore, μ^* falls as horizon lengthens. For $H = T$, the unbiased forecast compounds initial portfolio value at the geometric average return. For even longer horizons, one would apply a weight greater than 1.0 to the geometric average, and a negative weight to the arithmetic average, resulting in a growth rate below either the geometric or arithmetic mean.[7]

This analysis sheds light on an apparent paradox. Suppose we assume that returns come from a distribution that is stable over time. In that case, the 76-year historic return from the *SBBI* data base ending in 2001 would be a reasonable (albeit imprecise) estimate of cumulative return over the next 76 years. Compounding at the historic *geometric* average over a 76-year horizon would (by construction) match the proportional growth in wealth

---

[6] This result is similar to that of Blume (1974). He shows that final wealth is *approximately* a weighted average of wealth relatives based on geometric and arithmetic means. In contrast, we are able to obtain an exact result by focusing on drift rates of wealth (rather than wealth per se) under log-normality.

[7] Cooper (1996) analyzes bias due to estimation error in the context of discount factors. However, because discount factors involve powers of the reciprocal of the rate of return, the biases he finds differ from those here. He finds that both arithmetic and geometric averages result in downward biased estimators of the appropriate discount factor, but that the arithmetic average is typically very close to the unbiased discount rate. In contrast, we find that the arithmetic average always results in upward bias, that the bias can be substantial, and that the geometric average can result in either upward or downward bias.

14

realized over the past 76 years. In contrast, compounding at the sample arithmetic average for 76 years (as typically prescribed by the literature) would necessarily give a forecast of growth in wealth *greater* than the one realized historically. Following standard practice thus *ensures* a forecast of future portfolio growth that exceeds historic experience. The bias correction described in equation (6) shows that this forecasting exercise (with $T = H = 76$ years) actually would call for compounding at the geometric average, so that the forecast of 76-year cumulative return would match that experienced historically.

2. Indicative biases

Is the potential bias in forecasts of cumulative returns economically significant? Unfortunately, it seems to be.

Assume that $\hat{\mu} = .10$ and $\sigma = .20$. Figure 2, Panel A, shows the forecasted growth of funds over investment horizons ranging from 0 to 40 years for four forecasts: arithmetic average (labeled A), geometric average (labeled G), and two unbiased growth rates (labeled U), computed based on historical sample periods of different durations. (The unbiased estimator weights the geometric and arithmetic rates using weights H/T and $1 - H/T$, so each sample period results in a different estimator.) We assume in Figure 2 that $\hat{\mu}$ is estimated using either a 76-year sample period (corresponding to the *SBBI* period for the U.S.) or a 30-year period, a sample period that might be available for an emerging market, or even for a developed capital market such as that of U.S. if one believes that the post-Vietnam era represents a structural economic break.

[Figure 2 here]

15

Panel A shows that the arithmetic estimator suffers from severe upward bias at long horizons. In contrast, the bias of the geometric estimator depends on the relation between H and T. When they are close, the estimator is relatively unbiased. In fact, the Figure illustrates that for $H = T = 30$ years, the geometric and unbiased estimators are equal. In general in Panel A, the geometric and unbiased estimators for $T = 30$ years do not diverge much for investment horizons less than 35 years. However, when the discrepancy between H and T is greater, for example, when $T = 76$ years, the downward bias in the geometric estimator can be profound, in fact, roughly equal to the upward bias in the arithmetic estimator at the equivalent horizon.

Panel B presents another view of the relative biases. It presents the *ratios* of arithmetic or geometric forecasts of cumulative return to unbiased forecasts. For $T = 30$ years, the bias in the arithmetic estimator rises dramatically with investment horizon: at a horizon of $H = 20$ years, bias is about 30 percent, but at a 40-year horizon, bias rises to almost 200 percent. For $T = 76$ years, the arithmetic estimator performs much better, but still is subject to an upward bias of about 50 percent at a horizon of 40 years. Using longer sample periods, T, the arithmetic forecast obviously performs better; however, the stability of the underlying return process at ever-longer horizons becomes increasingly suspect. Symmetrically, for long T, the geometric estimator can be severely downward biased. Using $T = 76$ years of data and an investment horizon of $H = 40$ years, for example, the geometric forecast of final wealth is only about 60 percent of the unbiased forecast.

The tradeoff between long sample periods (which would increase precision when the underlying return process is stable) versus truncated sample periods (which disregard old and possibly no-longer-representative experience) is highlighted in Table 2, which presents

16

40-year return forecasts for a small sample of countries based on historical sample periods of different lengths. The longest series available from Datastream for France, Germany, and the U.K. are well more than a century in length. For the U.K, with a 201-year data span, the arithmetic average return is almost equal to the unbiased compounding rate even for a horizon as long as 40 years. But notice for these countries that the estimates of μ over these long periods are far lower than estimates derived over the last 82 years, and smaller still than estimates based on the latest 52 years. Do we conclude that the estimates based on longer data series are more reliable by virtue of their higher sample size, or that structural change over the last century or two makes returns from the nineteenth century of dubious value for predicting twenty-first century returns?

We would point out that *regardless* of how one might answer this question, arithmetic averages based on relatively recent experience will result in compounding rates that are far too high, at least for these countries. If we accept the long-duration estimates, the resulting values of $\hat{\mu}$ are far lower than estimates obtained from more recent periods. In contrast, if we rely on only more recent data we obtain higher $\hat{\mu}$, but now the unbiased estimator gives higher weight to the (lower) geometric average (since H/T is larger). In either case, the compound growth rates that give unbiased forecasts of cumulative returns are substantially below results using conventional arithmetic averages.

The impact of these considerations on forecasts of future portfolio value is dramatic. For any of these countries, compounding at the arithmetic average return calculated from sample periods of either the most recent 82 or 52 years results in forecasts of future value that are roughly double the corresponding unbiased forecasts based on the same data periods.

17

Emerging markets present an even greater problem. At the bottom of Table 2, we present estimates for Hong Kong and the MSCI emerging market index (both in dollar-denominated returns), where available data span only 17.6 or 13.6 years respectively, and volatility is high (especially for Hong Kong). With such short historical estimation periods, the unbiased forecast of future value at $H = 40$ is below *both* the geometric and arithmetic forecasts, and by enormous margins. For the emerging market index, with T equal to only 13.6 years, the unbiased forecast is only 3.1% the forecast obtained by compounding at the arithmetic average, and is only 10.1% of the compound geometric average.

3. Conclusion

A longstanding debate on forecasting future portfolio value has focused on the relative merits of the geometric versus arithmetic average return as a compounding rate. We have shown in this paper that when these averages must be estimated subject to sampling error, neither approach yields unbiased forecasts. For typical investment horizons, the proper compounding rate is in-between these two values. Specifically, unbiased estimates of future portfolio value require that the current value be compounded forward at a weighted average of the arithmetic and geometric rates. The proper weight placed on the geometric rate equals the ratio of the investment horizon to the sample estimation period. Therefore, for short investment horizons, the arithmetic average will be close to the "unbiased compounding rate." However, as the horizon approaches the length of the estimation period, the weight on the geometric average approaches 1. For even longer horizons, both the geometric and arithmetic average forecasts will be upward biased.

18

We demonstrate that these biases can be empirically significant. For investment horizons of 40 years, for example, the difference in forecasts of cumulative performance easily can exceed a factor of two. The percentage differences in forecasts grows along with the investment horizon, as well as with the imprecision in the estimate of the mean return.

For future research we reserve questions concerning the robustness of our results to distributional assumptions concerning rates of return. From preliminary research (Jacquier, Kane and Marcus, 2002) the results appear robust to heteroskedasticity and serial correlation. We also address there forecasting criteria other than unbiasedness, for example, minimum mean squared error forecasts. These forecasts also depend critically on the ratio of the investment horizon to the sample estimation period, and present qualitatively similar patterns as those presented in this paper.

References

Blume, M. 1974. "Unbiased Estimates of Long-Run Expected Rates of Return," *Journal of the American Statistical Association* vol. 69, no. 347 (September): 634-638.

Bodie, Z., Kane A., and Marcus, A.J. 2002. *Investments*, 5th edition. New York: McGraw-Hill Irwin.

Brealey, R. and Myers, S.C. 2003. *Principles of Corporate Finance*, 7th edition. New York: McGraw-Hill Irwin.

Cooper, I. 1996. "Arithmetic Versus Geometric Mean Estimators: Setting Discount Rates for Capital Budgeting," *European Financial Management* vol. 2, no. 2 (July): 157-167.

Fama, E. and French, K. 2002. "The Equity Premium," *Journal of Finance* vol. 57 (April): 637-660.

Jacquier, E., Kane, A. and Marcus, A.J. 2002. "Optimal Forecasts of Long-Term Returns and Asset Allocation: Geometric, Arithmetic, or Other Means?" Working paper.

Jagannathan, R., McGrattan, E.R. and Scherbina, A. 2000. "The Declining U.S. Equity Premium," *Federal Reserve Bank of Minneapolis Quarterly Review* vol. 24 (Fall): 3-19.

Merton, R. 1980. "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics* 8 (December): 323-406.

Ross, S.A. Westerfield and R.W. Jaffe, J. 2002. *Corporate Finance*, 6th edition, New York: McGraw-Hill Irwin.

Table 1: Bias induced by forecasting final portfolio value using arithmetic average return of portfolio over a sample period.   Ratio of forecast to true expected value of cumulative return. [σ is the standard deviation of annual returns.]

A.　　　　　**Sample period = 76 years**

	Horizon (years)			
σ	10	20	30	40
0.15	1.015	1.062	1.145	1.271
0.2	1.027	1.113	1.271	1.532
0.25	1.043	1.181	1.455	1.948
0.3	1.062	1.271	1.716	2.612

B.　　　　　**Sample period = 30 years**

	Horizon (years)			
σ	10	20	30	40
0.15	1.038	1.162	1.401	1.822
0.2	1.069	1.306	1.822	2.906
0.25	1.110	1.517	2.554	5.294
0.3	1.162	1.822	3.857	11.023

Table 2: Estimates of compounding rates and future portfolio values. A, G, U: are annual compounding rates using arithmetic average A, geometric average G, and unbiased estimator U (H = 40) computed over the sample period specified. Corresponding forecasts of future portfolio values (relative to an initial $1 investment) for an investment horizon H = 40 years are denoted V(A), V(G), V(U).

Country/Index	Freq[a]	T (years)	Begin	End	Sample estimates $\hat{\mu}$	$\hat{\sigma}$	Compound growth rates A	G	U	Future portfolio value V(A)	V(G)	V(U)
Canada/TSE	A	78	1914	2001	4.8	16.7	6.4	4.9	5.6	11.9	6.8	9.0
Canada/TSE	A	52	1950	2001	6.6	14.9	8.0	6.8	7.1	21.8	14.0	15.5
France/SBF250	A	145	1857	2001	5.1	19.7	7.3	5.2	6.7	16.7	7.7	13.5
France/SBF250	A	82	1920	2001	8.5	24.7	12.2	8.9	10.6	101.5	30.0	56.0
France/SBF250	A	52	1950	2001	8.7	22.2	11.8	9.1	9.7	87.0	32.5	40.8
Germany/DAX	A	145	1857	2001	1.9	32.2	7.3	1.9	5.8	17.0	2.1	9.6
Germany/DAX	A	82	1920	2001	5.5	37.0	13.1	5.7	9.4	139.5	9.0	36.7
Germany/DAX	A	52	1950	2001	8.0	22.8	11.2	8.3	9.0	69.4	24.5	31.2
UK/FTAS	A	201	1801	2001	2.4	15.6	3.7	2.4	3.4	4.2	2.6	3.9
UK/FTAS	A	82	1920	2001	5.5	20.0	7.8	5.7	6.7	20.1	9.0	13.6
UK/FTAS	A	52	1950	2001	6.4	24.7	9.9	6.6	8.3	43.8	12.9	24.2
Japan/Nikkei	A	52	1950	2001	8.8	24.1	12.4	9.2	9.9	107.9	33.8	44.2
Hong Kong[b]	D	27.6	5/31/73	1/2/01	10.7	30.7	16.7	11.3	9.0	475.8	72.2	31.0
MSCI/$Emg Mkt	M	13.6	1/88	7/01	8.2	24.2	11.8	8.5	2.5	85.7	26.6	2.7

a: Frequency: A annual; M monthly; W weekly; D daily
b: $HangSeng Without dividend yield. mu=13.5 with dividend yield

Figure 1: Probability densities of estimates of $\hat{\mu}$, the unbiased estimate of μ, and estimates of final portfolio value, $e^{(\hat{\mu} + \frac{1}{2}\sigma^2)H}$. $\hat{\mu}$ is symmetrically distributed around μ (Panel A), but the distributions of estimates of final wealth obtained by compounding at the arithmetic average rate of return are skewed. Skewness and bias of the estimate of final value increase with the investment horizon, H. Panels B, C, and D show investment horizons $H = 2$, 10, and 20 years respectively. Annual returns are assumed to be log-normal with parameters $\mu = .10$ and $\sigma = .20$.
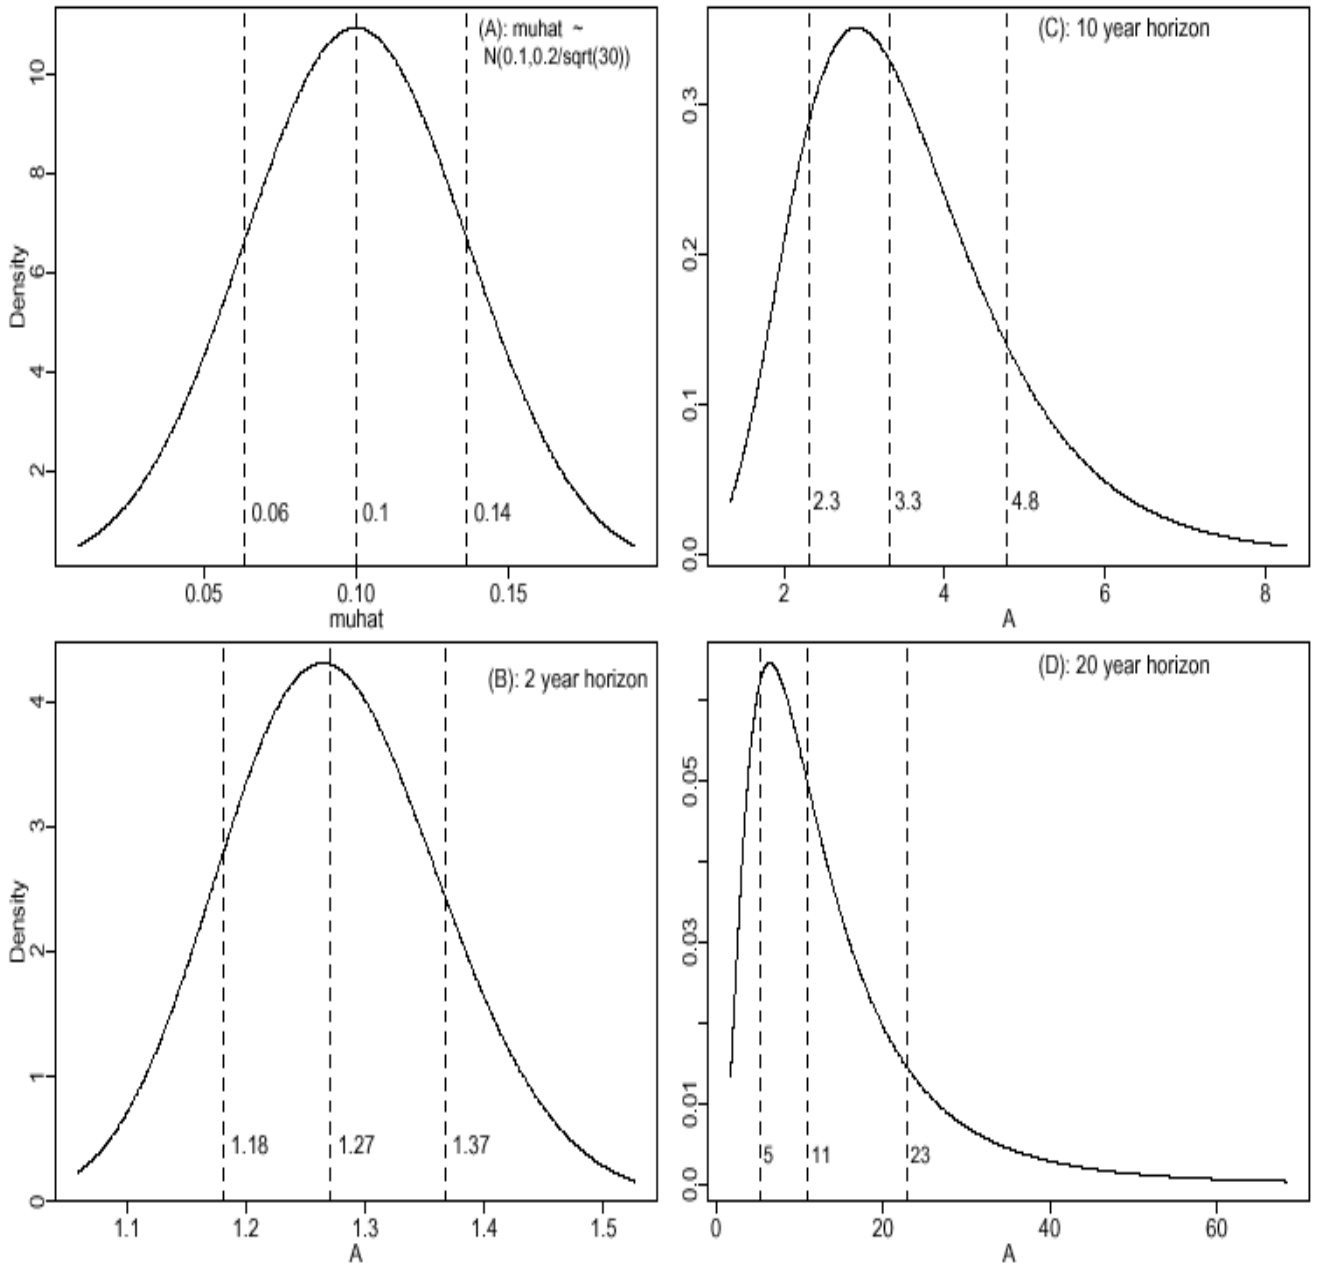
Figure 2: Forecasts of final portfolio value. Initial investment = $1. Annual returns are assumed to be log-normal with parameters $\hat{\mu} = .10$ and $\sigma = .20$. Forecast of final value using arithmetic average return is $A = e^{(\hat{\mu} + \frac{1}{2}\sigma^2)H}$, using arithmetic average return is $G = e^{\hat{\mu} H}$. The unbiased forecast is $U = e^{(\mu^* + \frac{1}{2}\sigma^2)H}$ which amounts to compounding forward at a rate equal to a weighted average of the geometric and arithmetic averages with weights H/T and $1 - H/T$. Panel (a) presents forecasts of final value Panel (b) presents ratios of the arithmetic or geometric forecast to the unbiased forecast.