
The Statistics of Sharpe Ratios

Andrew W. Lo

The building blocks of the Sharpe ratio—expected returns and volatilities—are unknown quantities that must be estimated statistically and are, therefore, subject to estimation error. This raises the natural question: How accurately are Sharpe ratios measured? To address this question, I derive explicit expressions for the statistical distribution of the Sharpe ratio using standard asymptotic theory under several sets of assumptions for the return-generating process—independently and identically distributed returns, stationary returns, and with time aggregation. I show that monthly Sharpe ratios cannot be annualized by multiplying by $\sqrt{12}$ except under very special circumstances, and I derive the correct method of conversion in the general case of stationary returns. In an illustrative empirical example of mutual funds and hedge funds, I find that the annual Sharpe ratio for a hedge fund can be overstated by as much as 65 percent because of the presence of serial correlation in monthly returns, and once this serial correlation is properly taken into account, the rankings of hedge funds based on Sharpe ratios can change dramatically.

One of the most commonly cited statistics in financial analysis is the Sharpe ratio, the ratio of the excess expected return of an investment to its return volatility or standard deviation. Originally motivated by mean-variance analysis and the Sharpe-Lintner Capital Asset Pricing Model, the Sharpe ratio is now used in many different contexts, from performance attribution to tests of market efficiency to risk management.¹ Given the Sharpe ratio's widespread use and the myriad interpretations that it has acquired over the years, it is surprising that so little attention has been paid to its statistical properties. Because expected returns and volatilities are quantities that are generally not observable, they must be estimated in some fashion. The inevitable estimation errors that arise imply that the Sharpe ratio is also estimated with error, raising the natural question: How accurately are Sharpe ratios measured?

In this article, I provide an answer by deriving the statistical distribution of the Sharpe ratio using standard econometric methods under several different sets of assumptions for the statistical behavior of the return series on which the Sharpe ratio is based. Armed with this statistical distribution, I

show that confidence intervals, standard errors, and hypothesis tests can be computed for the estimated Sharpe ratio in much the same way that they are computed for regression coefficients such as portfolio alphas and betas.

The accuracy of Sharpe ratio estimators hinges on the statistical properties of returns, and these properties can vary considerably among portfolios, strategies, and over time. In other words, the Sharpe ratio estimator's statistical properties typically will depend on the investment style of the portfolio being evaluated. At a superficial level, the intuition for this claim is obvious: The performance of more volatile investment strategies is more difficult to gauge than that of less volatile strategies. Therefore, it should come as no surprise that the results derived in this article imply that, for example, Sharpe ratios are likely to be more accurately estimated for mutual funds than for hedge funds.

A less intuitive implication is that the time-series properties of investment strategies (e.g., mean reversion, momentum, and other forms of serial correlation) can have a nontrivial impact on the Sharpe ratio estimator itself, especially in computing an annualized Sharpe ratio from monthly data. In particular, the results derived in this article show that the common practice of annualizing Sharpe ratios by multiplying monthly estimates by $\sqrt{12}$ is correct only under very special circumstances and that the correct multiplier—which depends on the serial correlation of the portfolio's

Andrew W. Lo is Harris & Harris Group Professor at the Sloan School of Management, Massachusetts Institute of Technology, Cambridge, and chief scientific officer for AlphaSimplex Group, LLC, Cambridge.

returns—can yield Sharpe ratios that are considerably smaller (in the case of positive serial correlation) or larger (in the case of negative serial correlation). Therefore, Sharpe ratio estimators must be computed and interpreted in the context of the particular investment style with which a portfolio's returns have been generated.

Let R_t denote the one-period simple return of a portfolio or fund between dates $t - 1$ and t and denote by μ and σ^2 its mean and variance:

$$\mu \equiv E(R_t), \quad (1a)$$

and

$$\sigma^2 \equiv \text{Var}(R_t). \quad (1b)$$

Recall that the Sharpe ratio (SR) is defined as the ratio of the excess expected return to the standard deviation of return:

$$\text{SR} \equiv \frac{\mu - R_f}{\sigma}, \quad (2)$$

where the excess expected return is usually computed relative to the risk-free rate, R_f . Because μ and σ are the population moments of the distribution of R_t , they are unobservable and must be estimated using historical data.

Given a sample of historical returns (R_1, R_2, \dots, R_T), the standard estimators for these moments are the sample mean and variance:

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T R_t \quad (3a)$$

and

$$\hat{\sigma}^2 = \frac{1}{T} \sum_{t=1}^T (R_t - \hat{\mu})^2, \quad (3b)$$

from which the estimator of the Sharpe ratio ($\widehat{\text{SR}}$) follows immediately:

$$\widehat{\text{SR}} = \frac{\hat{\mu} - R_f}{\hat{\sigma}}. \quad (4)$$

Using a set of techniques collectively known as “large-sample” or “asymptotic” statistical theory in which the Central Limit Theorem is applied to estimators such as $\hat{\mu}$ and $\hat{\sigma}^2$, the distribution of $\widehat{\text{SR}}$ and other nonlinear functions of $\hat{\mu}$ and $\hat{\sigma}^2$ can be easily derived.

In the next section, I present the statistical distribution of $\widehat{\text{SR}}$ under the standard assumption that returns are independently and identically distributed (IID). This distribution completely characterizes the statistical behavior of $\widehat{\text{SR}}$ in large samples and allows us to quantify the precision with which $\widehat{\text{SR}}$ estimates SR. But because the IID assumption is extremely restrictive and often violated by financial

data, a more general distribution is derived in the “Non-IID Returns” section, one that applies to returns with serial correlation, time-varying conditional volatilities, and many other characteristics of historical financial time series. In the “Time Aggregation” section, I develop explicit expressions for “time-aggregated” Sharpe ratio estimators (e.g., expressions for converting monthly Sharpe ratio estimates to annual estimates) and their distributions. To illustrate the practical relevance of these estimators, I apply them to a sample of monthly mutual fund and hedge fund returns and show that serial correlation has dramatic effects on the annual Sharpe ratios of hedge funds, inflating Sharpe ratios by more than 65 percent in some cases and deflating Sharpe ratios in other cases.

IID Returns

To derive a measure of the uncertainty surrounding the estimator $\widehat{\text{SR}}$, we need to specify the statistical properties of R_t because these properties determine the uncertainty surrounding the component estimators $\hat{\mu}$ and $\hat{\sigma}^2$. Although this may seem like a theoretical exercise best left for statisticians—not unlike the specification of the assumptions needed to yield well-behaved estimates from a linear regression—there is often a direct connection between the investment management process of a portfolio and its statistical properties. For example, a change in the portfolio manager's style from a small-cap value orientation to a large-cap growth orientation will typically have an impact on the portfolio's volatility, degree of mean reversion, and market beta. Even for a fixed investment style, a portfolio's characteristics can change over time because of fund inflows and outflows, capacity constraints (e.g., a microcap fund that is close to its market-capitalization limit), liquidity constraints (e.g., an emerging market or private equity fund), and changes in market conditions (e.g., sudden increases or decreases in volatility, shifts in central banking policy, and extraordinary events, such as the default of Russian government bonds in August 1998). Therefore, the investment style and market environment must be kept in mind when formulating the assumptions for the statistical properties of a portfolio's returns.

Perhaps the simplest set of assumptions that we can specify for R_t is that they are independently and identically distributed. This means that the probability distribution of R_t is identical to that of R_s for any two dates t and s and that R_t and R_s are statistically independent for all $t \neq s$. Although these conditions are extreme and empirically implausible—the probability distribution of the monthly return of the S&P 500 Index in October 1987 is likely to differ

from the probability distribution of the monthly return of the S&P 500 in December 2000—they provide an excellent starting point for understanding the statistical properties of Sharpe ratios. In the next section, these assumptions will be replaced with a more general set of conditions for returns.

Under the assumption that returns are IID and have finite mean μ and variance σ^2 , it is well known that the estimators $\hat{\mu}$ and $\hat{\sigma}^2$ in Equation 3 have the following normal distributions in large samples, or “asymptotically,” due to the Central Limit Theorem:²

$$\sqrt{T}(\hat{\mu} - \mu) \stackrel{a}{\sim} N(0, \sigma^2), \sqrt{T}(\hat{\sigma}^2 - \sigma^2) \stackrel{a}{\sim} N(0, 2\sigma^4), \quad (5)$$

where $\stackrel{a}{\sim}$ denotes the fact that this relationship is an asymptotic one [i.e., as T increases without bound, the probability distributions of $\sqrt{T}(\hat{\mu} - \mu)$ and $\sqrt{T}(\hat{\sigma}^2 - \sigma^2)$ approach the normal distribution, with mean zero and variances σ^2 and $2\sigma^4$, respectively]. These asymptotic distributions imply that the estimation error of $\hat{\mu}$ and $\hat{\sigma}^2$ can be approximated by

$$\text{Var}(\hat{\mu}) \stackrel{a}{=} \frac{\sigma^2}{T}, \text{Var}(\hat{\sigma}^2) \stackrel{a}{=} \frac{2\sigma^4}{T}, \quad (6)$$

where $\stackrel{a}{=}$ indicates that these relations are based on asymptotic approximations. Note that in Equation 6, the variances of both estimators approach zero as T increases, reflecting the fact that the estimation errors become smaller as the sample size grows. An additional property of $\hat{\mu}$ and $\hat{\sigma}$ in the special case of IID returns is that they are statistically independent in large samples, which greatly simplifies our analysis of the statistical properties of functions of these estimators (e.g., the Sharpe ratio).

Now, denote by the function $g(\mu, \sigma^2)$ the Sharpe ratio defined in Equation 2; hence, the Sharpe ratio estimator is simply $g(\hat{\mu}, \hat{\sigma}^2) = \widehat{\text{SR}}$. When the Sharpe ratio is expressed in this form, it is apparent that the estimation errors in $\hat{\mu}$ and $\hat{\sigma}^2$ will affect $g(\hat{\mu}, \hat{\sigma}^2)$ and that the nature of these effects depends critically on the properties of the function g . Specifically, in the “IID Returns” section of Appendix A, I show that the asymptotic distribution of the Sharpe ratio estimator is $\sqrt{T}(\widehat{\text{SR}} - \text{SR}) \stackrel{a}{\sim} N(0, V_{\text{IID}})$, where the asymptotic variance is given by the following weighted average of the asymptotic variances of $\hat{\mu}$ and $\hat{\sigma}^2$:

$$V_{\text{IID}} = \left(\frac{\partial g}{\partial \mu}\right)^2 \sigma^2 + \left(\frac{\partial g}{\partial \sigma^2}\right)^2 2\sigma^4. \quad (7)$$

The weights in Equation 7 are simply the squared sensitivities of g with respect to μ and σ^2 , respectively: The more sensitive g is to a particular parameter, the more influential its asymptotic variance will be in the weighted average. This relationship

is reminiscent of the expression for the variance of the weighted sum of two random variables, except that in Equation 7, there is no covariance term. This is due to the fact that $\hat{\mu}$ and $\hat{\sigma}^2$ are asymptotically independent, thanks to our simplifying assumption of IID returns. In the next sections, the IID assumption will be replaced by a more general set of conditions on returns, in which case, the covariance between $\hat{\mu}$ and $\hat{\sigma}^2$ will no longer be zero and the corresponding expression for the asymptotic variance of the Sharpe ratio estimator will be somewhat more involved.

The asymptotic variance of $\widehat{\text{SR}}$ given in Equation 7 can be further simplified by evaluating the sensitivities explicitly— $\partial g/\partial \mu = 1/\sigma$ and $\partial g/\partial \sigma^2 = -(\mu - R_f)/(2\sigma^3)$ —and then combining terms to yield

$$\begin{aligned} V_{\text{IID}} &= 1 + \frac{(\mu - R_f)^2}{2\sigma^2} \\ &= 1 + \frac{1}{2}\text{SR}^2. \end{aligned} \quad (8)$$

Therefore, standard errors (SEs) for the Sharpe ratio estimator $\widehat{\text{SR}}$ can be computed as

$$\text{SE}(\widehat{\text{SR}}) \stackrel{a}{=} \sqrt{\left(1 + \frac{1}{2}\text{SR}^2\right)/T}, \quad (9)$$

and this quantity can be estimated by substituting $\widehat{\text{SR}}$ for SR. Confidence intervals for SR can also be constructed from Equation 9; for example, the 95 percent confidence interval for SR around the estimator $\widehat{\text{SR}}$ is simply

$$\widehat{\text{SR}} \pm 1.96 \times \sqrt{\left(1 + \frac{1}{2}\widehat{\text{SR}}^2\right)/T}. \quad (10)$$

Table 1 reports values of Equation 9 for various combinations of Sharpe ratios and sample sizes. Observe that for any given sample size T , larger Sharpe ratios imply larger standard errors. For example, in a sample of 60 observations, the standard error of the Sharpe ratio estimator is 0.188 when the true Sharpe ratio is 1.50 but is 0.303 when the true Sharpe ratio is 3.00. This implies that the performance of investments such as hedge funds, for which high Sharpe ratios are one of the primary objectives, will tend to be less precisely estimated. However, as a percentage of the Sharpe ratio, the standard error given by Equation 9 does approach a finite limit as SR increases since

$$\frac{\text{SE}(\widehat{\text{SR}})}{\widehat{\text{SR}}} = \sqrt{\frac{1 + (1/2)\text{SR}^2}{T\text{SR}^2}} \rightarrow \frac{\sqrt{1}}{\sqrt{2T}} \quad (11)$$

as SR increases without bound. Therefore, the uncertainty surrounding the IID Sharpe ratio estimator will be approximately the same proportion

Table 1. Asymptotic Standard Errors of Sharpe Ratio Estimators for Combinations of Sharpe Ratio and Sample Size

SR	Sample Size, T							
	12	24	36	48	60	125	250	500
0.50	0.306	0.217	0.177	0.153	0.137	0.095	0.067	0.047
0.75	0.327	0.231	0.189	0.163	0.146	0.101	0.072	0.051
1.00	0.354	0.250	0.204	0.177	0.158	0.110	0.077	0.055
1.25	0.385	0.272	0.222	0.193	0.172	0.119	0.084	0.060
1.50	0.421	0.298	0.243	0.210	0.188	0.130	0.092	0.065
1.75	0.459	0.325	0.265	0.230	0.205	0.142	0.101	0.071
2.00	0.500	0.354	0.289	0.250	0.224	0.155	0.110	0.077
2.25	0.542	0.384	0.313	0.271	0.243	0.168	0.119	0.084
2.50	0.586	0.415	0.339	0.293	0.262	0.182	0.128	0.091
2.75	0.631	0.446	0.364	0.316	0.282	0.196	0.138	0.098
3.00	0.677	0.479	0.391	0.339	0.303	0.210	0.148	0.105

Note: Returns are assumed to be IID, which implies $V_{IID} = 1 + 1/2SR^2$.

of the Sharpe ratio for higher Sharpe ratio investments with the same number of observations T .

We can develop further intuition for the impact of estimation errors in $\hat{\mu}$ and $\hat{\sigma}^2$ on the Sharpe ratio by calculating the proportion of asymptotic variance that is attributable to $\hat{\mu}$ versus $\hat{\sigma}^2$. From Equation 7, the fraction of V_{IID} due to estimation error in $\hat{\mu}$ versus $\hat{\sigma}^2$ is simply

$$\frac{(\partial g / \partial \mu)^2 \sigma^2}{V_{IID}} = \frac{1}{1 + (1/2)SR^2} \quad (12a)$$

and

$$\frac{(\partial g / \partial \sigma)^2 2\sigma^4}{V_{IID}} = \frac{(1/2)SR^2}{1 + (1/2)SR^2}. \quad (12b)$$

For a small Sharpe ratio, such as 0.25, this proportion—which depends only on the true Sharpe ratio—is 97.0 percent, indicating that most of the variability in the Sharpe ratio estimator is a result of variability in $\hat{\mu}$. However, for higher Sharpe ratios, the reverse is true: For $SR = 2.00$, only 33.3 percent of the variability of \widehat{SR} comes from $\hat{\mu}$, and for a Sharpe ratio of 3.00, only 18.2 percent of the estimator error of \widehat{SR} is attributable to $\hat{\mu}$.

Non-IID Returns

Many studies have documented various violations of the assumption of IID returns for financial securities;³ hence, the results of the previous section may be of limited practical value in certain circumstances. Fortunately, it is possible to derive similar results under more general conditions, conditions that allow for serial correlation, conditional heteroskedasticity, and other forms of dependence and heterogeneity in returns. In particular, if

returns satisfy the assumption of “stationarity,” then a version of the Central Limit Theorem still applies to most estimators and the corresponding asymptotic distribution can be derived. The formal definition of stationarity is that the joint probability distribution $F(R_{t_1}, R_{t_2}, \dots, R_{t_n})$ of an arbitrary collection of returns $R_{t_1}, R_{t_2}, \dots, R_{t_n}$ does not change if all the dates are incremented by the same number of periods; that is,

$$F(R_{t_1+k}, R_{t_2+k}, \dots, R_{t_n+k}) = F(R_{t_1}, R_{t_2}, \dots, R_{t_n}) \quad (13)$$

for all k . Such a condition implies that mean μ and variance σ^2 (and all higher moments) are constant over time but otherwise allows for quite a broad set of dynamics for R_t , including serial correlation, dependence on such factors as the market portfolio, time-varying conditional volatilities, jumps, and other empirically relevant phenomena.

Under the assumption of stationarity,⁴ a version of the Central Limit Theorem can still be applied to the estimator \widehat{SR} . However, in this case, the expression for the variance of \widehat{SR} is somewhat more complex because of the possibility of dependence between the components $\hat{\mu}$ and $\hat{\sigma}^2$. In the “Non-IID Returns” section of Appendix A, I show that the asymptotic distribution can be derived by using a “robust” estimator—an estimator that is effective under many different sets of assumptions for the statistical properties of returns—to estimate the Sharpe ratio.⁵ In particular, I use a *generalized method of moments* (GMM) estimator to estimate $\hat{\mu}$ and $\hat{\sigma}^2$, and the results of Hansen (1982) can be used to obtain the following asymptotic distribution:

$$\sqrt{T}(\widehat{SR} - SR) \stackrel{d}{\sim} N(0, V_{GMM}), V_{GMM} = \frac{\partial g}{\partial \theta} \Sigma \frac{\partial g}{\partial \theta}, \quad (14)$$

where the definitions of $\partial g/\partial \theta$ and Σ and a method for estimating them are given in the second section of Appendix A. Therefore, for non-IID returns, the standard error of the Sharpe ratio can be estimated by

$$SE[\widehat{SR}] \stackrel{a}{=} \sqrt{\widehat{V}_{GMM}}/T \tag{15}$$

and confidence intervals for SR can be constructed in a similar fashion to Equation 10.

Time Aggregation

In many applications, it is necessary to convert Sharpe ratio estimates from one frequency to another. For example, a Sharpe ratio estimated from monthly data cannot be directly compared with one estimated from annual data; hence, one statistic must be converted to the same frequency as the other to yield a fair comparison. Moreover, in some cases, it is possible to derive a more precise estimator of an annual quantity by using monthly or daily data and then performing time aggregation instead of estimating the quantity directly using annual data.⁶

In the case of Sharpe ratios, the most common method for performing such time aggregation is to multiply the higher-frequency Sharpe ratio by the square root of the number of periods contained in the lower-frequency holding period (e.g., multiply a monthly estimator by $\sqrt{12}$ to obtain an annual estimator). In this section, I show that this rule of thumb is correct only under the assumption of IID returns. For non-IID returns, an alternative procedure must be used, one that accounts for serial correlation in returns in a very specific manner.

IID Returns. Consider first the case of IID returns. Denote by $R_t(q)$ the following q -period return:

$$R_t(q) \equiv R_t + R_{t-1} + \dots + R_{t-q+1}, \tag{16}$$

where I have ignored the effects of compounding for computational convenience.⁷ Under the IID assumption, the variance of $R_t(q)$ is directly proportional to q ; hence, the Sharpe ratio satisfies the simple relationship:

$$\begin{aligned} SR(q) &= \frac{E[R_t(q)] - R_f(q)}{\sqrt{\text{Var}[R_t(q)]}} \\ &= \frac{q(\mu - R_f)}{\sqrt{q}\sigma} \\ &= \sqrt{q}SR. \end{aligned} \tag{17}$$

Despite the fact that the Sharpe ratio may seem to be “unitless” because it is the ratio of two quantities with the same units, it does depend on the timescale with respect to which the numerator and denominator are defined. The reason is that the

numerator increases linearly with aggregation value q whereas the denominator increases as the square root of q under IID returns; hence, the ratio will increase as the square root of q , making a longer-horizon investment seem more attractive. This interpretation is highly misleading and should not be taken at face value. Indeed, the Sharpe ratio is not a complete summary of the risks of a multiperiod investment strategy and should never be used as the sole criterion for making an investment decision.⁸

The asymptotic distribution of $\widehat{SR}(q)$ follows directly from Equation 17 because $\widehat{SR}(q)$ is proportional to SR:

$$\begin{aligned} \sqrt{T}[\widehat{SR}(q) - \sqrt{q}SR] &\stackrel{a}{\sim} N\left[0, V_{IID}(q)\right], \\ V_{IID}(q) &= qV_{IID} = q\left(1 + \frac{1}{2}SR^2\right). \end{aligned} \tag{18}$$

Non-IID Returns. The relationship between SR and $SR(q)$ is somewhat more involved for non-IID returns because the variance of $R_t(q)$ is not just the sum of the variances of component returns but also includes all the covariances. Specifically, under the assumption that returns R_t are stationary,

$$\begin{aligned} \text{Var}[R_t(q)] &= \sum_{i=0}^{q-1} \sum_{j=0}^{q-1} \text{Cov}(R_{t-i}, R_{t-j}) \\ &= q\sigma^2 + 2\sigma^2 \sum_{k=1}^{q-1} (q-k)\rho_k, \end{aligned} \tag{19}$$

where $\rho_k \equiv \text{Cov}(R_t, R_{t-k})/\text{Var}(R_t)$ is the k th-order autocorrelation of R_t .⁹ This yields the following relationship between SR and $SR(q)$:

$$SR(q) = \eta(q)SR, \eta(q) \equiv \frac{q}{\sqrt{q + 2\sum_{k=1}^{q-1} (q-k)\rho_k}}, \tag{20}$$

Note that Equation 20 reduces to Equation 17 if all autocorrelations ρ_k are zero, as in the case of IID returns. However, for non-IID returns, the adjustment factor for time-aggregated Sharpe ratios is generally not \sqrt{q} but a more complicated function of the first $q - 1$ autocorrelations of returns.

Example: First-Order Autoregressive Returns. To develop some intuition for the potential impact of serial correlation on the Sharpe ratio, consider the case in which returns follow a first-order autoregressive process or “AR(1)”:

$$R_t = \mu + \rho(R_{t-1} - \mu) + \varepsilon_t, \quad -1 < \rho < 1, \tag{21}$$

where ε_t is IID with mean zero and variance σ_ε^2 . In this case, the return in period t can be forecasted to some degree by the return in period $t - 1$, and this “autoregression” leads to serial correlation at all lags. In particular, Equation 21 implies that the

k th-order autocorrelation coefficient is simply ρ^k ; hence, the scale factor in Equation 20 can be evaluated explicitly as

$$\eta(q) = \sqrt{q} \left[1 + \frac{2\rho}{1-\rho} \left(1 - \frac{1-\rho^q}{q(1-\rho)} \right) \right]^{-1/2}. \quad (22)$$

Table 2 presents values of $\eta(q)$ for various values of ρ and q ; the row corresponding to $\rho = 0$ percent is the IID case in which the scale factor is simply \sqrt{q} . Note that for each holding-period q , positive serial correlation reduces the scale factor below the IID value and negative serial correlation increases it. The reason is that positive serial correlation implies that the variance of multiperiod returns increases faster than holding-period q ; hence, the variance of $R_t(q)$ is more than q times the variance of R_t , yielding a larger denominator in the Sharpe ratio than the IID case. For returns with negative serial correlation, the opposite is true: The variance of $R_t(q)$ is less than q times the variance of R_t , yielding a smaller denominator in the Sharpe ratio than the IID case. For returns with significant serial correlation, this effect can be substantial. For example, the annual Sharpe ratio of a portfolio with a monthly first-order autocorrelation of -20 percent is 4.17 times the monthly Sharpe ratio, whereas the

scale factor is 3.46 in the IID case and 2.88 when the monthly first-order autocorrelation is 20 percent.

These patterns are summarized in **Figure 1**, in which $\eta(q)$ is plotted as a function of q for five values of ρ . The middle ($\rho = 0$) curve corresponds to the standard scale factor \sqrt{q} , which is the correct

Figure 1. Scale Factors of Time-Aggregated Sharpe Ratios When Returns Follow an AR(1) Process: For $\rho = -0.50, -0.25, 0, 0.25,$ and 0.50

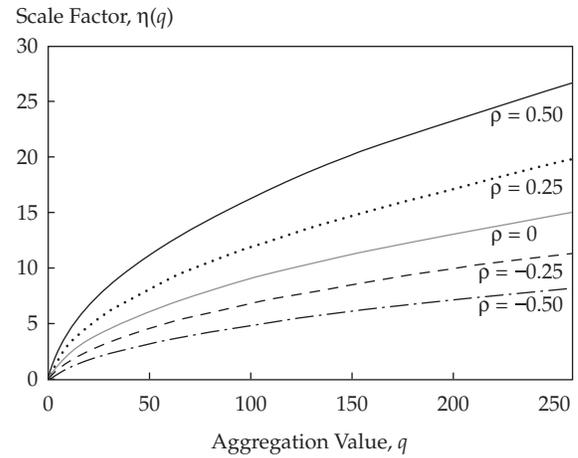


Table 2. Scale Factors for Time-Aggregated Sharpe Ratios When Returns Follow an AR(1) Process for Various Aggregation Values and First-Order Autocorrelations

ρ (%)	Aggregation Value, q									
	2	3	4	6	12	24	36	48	125	250
90	1.03	1.05	1.07	1.10	1.21	1.41	1.60	1.77	2.67	3.70
80	1.05	1.10	1.14	1.21	1.43	1.81	2.14	2.42	3.79	5.32
70	1.08	1.15	1.21	1.33	1.65	2.19	2.62	3.00	4.75	6.68
60	1.12	1.21	1.30	1.46	1.89	2.55	3.08	3.53	5.63	7.94
50	1.15	1.28	1.39	1.60	2.12	2.91	3.53	4.06	6.49	9.15
40	1.20	1.35	1.49	1.75	2.36	3.27	3.98	4.58	7.35	10.37
30	1.24	1.43	1.60	1.91	2.61	3.65	4.44	5.12	8.23	11.62
20	1.29	1.52	1.73	2.07	2.88	4.04	4.93	5.68	9.14	12.92
10	1.35	1.62	1.86	2.25	3.16	4.45	5.44	6.28	10.12	14.31
0	1.41	1.73	2.00	2.45	3.46	4.90	6.00	6.93	11.18	15.81
-10	1.49	1.85	2.16	2.66	3.80	5.39	6.61	7.64	12.35	17.47
-20	1.58	1.99	2.33	2.90	4.17	5.95	7.31	8.45	13.67	19.35
-30	1.69	2.13	2.53	3.17	4.60	6.59	8.10	9.38	15.20	21.52
-40	1.83	2.29	2.75	3.48	5.09	7.34	9.05	10.48	17.01	24.11
-50	2.00	2.45	3.02	3.84	5.69	8.26	10.21	11.84	19.26	27.31
-60	2.24	2.61	3.37	4.30	6.44	9.44	11.70	13.59	22.19	31.50
-70	2.58	2.76	3.86	4.92	7.45	11.05	13.77	16.04	26.33	37.43
-80	3.16	2.89	4.66	5.91	8.96	13.50	16.98	19.88	32.96	47.02
-90	4.47	2.97	6.47	8.09	12.06	18.29	23.32	27.61	46.99	67.65

factor when the correlation coefficient is zero. The curves above the middle one correspond to positive values of ρ , and those below the middle curve correspond to negative values of ρ . It is apparent that serial correlation has a nontrivial effect on the time aggregation of Sharpe ratios.

The General Case. More generally, using the expression for $\widehat{SR}(q)$ in Equation 20, we can construct an estimator of $SR(q)$ from estimators of the first $q - 1$ autocorrelations of R_t under the assumption of stationary returns. As in the “Non-IID Return” section, we can use GMM to estimate these autocorrelations as well as their asymptotic joint distribution, which can then be used to derive the following limiting distribution of $\widehat{SR}(q)$:

$$\sqrt{T}[\widehat{SR}(q) - SR(q)] \stackrel{d}{\rightarrow} N\left[0, V_{GMM}(q)\right], \tag{23}$$

$$V_{GMM}(q) = \frac{\partial g}{\partial \theta} \Sigma \frac{\partial g}{\partial \theta},$$

where the definitions of $\partial g / \partial \theta$ and Σ and formulas for estimating them are given in the “Time Aggregation” section of Appendix A. The standard error of $\widehat{SR}(q)$ is then given by

$$SE[\widehat{SR}(q)] \stackrel{d}{=} \sqrt{\widehat{V}_{GMM}(q) / T} \tag{24}$$

and confidence intervals can be constructed as in Equation 10.

Using $\widehat{V}_{GMM}(q)$ When Returns Are IID.

Although the robust estimator for $SR(q)$ is the appropriate estimator to use when returns are serially correlated or non-IID in other ways, there is a cost: additional estimation error induced by the autocovariance estimator, $\hat{\gamma}_k$, which manifests itself in the asymptotic variance, $\widehat{V}_{GMM}(q)$, of $\widehat{SR}(q)$. To develop a sense for the impact of estimation error on $\widehat{V}_{GMM}(q)$, consider the robust estimator when returns are, in fact, IID. In that case, $\gamma_k = 0$ for all $k > 0$ but because the robust estimator is a function of estimators $\hat{\gamma}_k$, the estimation errors of the autocovariance estimators will have an impact on $\widehat{V}_{GMM}(q)$. In particular, in the “Using $\widehat{V}_{GMM}(q)$ When Returns Are IID” section of Appendix A, I show that for IID returns, the asymptotic variance of robust estimator $\widehat{SR}(q)$ is given by

$$V_{GMM}(q) = \left[1 - \frac{v_3 SR}{\sigma^3} + (v_4 - \sigma^4) \frac{SR^2}{4\sigma^4} \right] + (\sqrt{q} SR)^2 \sum_{j=1}^{q-1} \left(1 - \frac{j}{q} \right)^2, \tag{25}$$

where $v_3 \equiv E[(R_t - \mu)^3]$ and $v_4 \equiv E[(R_t - \mu)^4]$ are the return’s third and fourth moments, respectively. Now suppose that returns are normally distributed. In that case, $v_3 = 0$ and $v_4 = 3\sigma^4$, which implies that

$$V_{GMM}(q) = V_{IID}(q) + (\sqrt{q} SR)^2 \sum_{j=1}^{q-1} \left(1 - \frac{j}{q} \right)^2 \geq V_{IID}(q). \tag{26}$$

The second term on the right side of Equation 26 represents the additional estimation error introduced by the estimated autocovariances in the more general estimator given in Equation A18 in Appendix A. By setting $q = 1$ so that no time aggregation is involved in the Sharpe ratio estimator (hence, no autocovariances enter into the estimator), the expression in Equation 26 reduces to the IID case given in Equation 18.

The asymptotic relative efficiency of $\widehat{SR}(q)$ can be evaluated explicitly by computing the ratio of $V_{GMM}(q)$ to $V_{IID}(q)$ in the case of IID normal returns:

$$\frac{V_{GMM}(q)}{V_{IID}(q)} = 1 + \frac{2 \sum_{j=1}^{q-1} (1 - j/q)^2}{1 + 2/SR^2}, \tag{27}$$

and **Table 3** reports these ratios for various combinations of Sharpe ratios and aggregation values q . Even for small aggregation values, such as $q = 2$, asymptotic variance $V_{GMM}(q)$ is significantly higher than $V_{IID}(q)$ —for example, 33 percent higher for a Sharpe ratio of 2.00. As the aggregation value increases, the asymptotic relative efficiency becomes even worse as more estimation error is built into the time-aggregated Sharpe ratio estimator. Even with a monthly Sharpe ratio of only 1.00, the annualized ($q = 12$) robust Sharpe ratio estimator has an asymptotic variance that is 334 percent of $V_{IID}(q)$.

The values in Table 3 suggest that, unless there is significant serial correlation in return series R_t , the robust Sharpe ratio estimator should not be used. A useful diagnostic to check for the presence of serial correlation is the Ljung–Box (1978) Q -statistic:

$$Q_{q-1} = T(T+2) \sum_{k=1}^{q-1} \frac{\hat{\rho}_k^2}{T-k}, \tag{28}$$

which is asymptotically distributed as χ_{q-1}^2 under the null hypothesis of no serial correlation.¹⁰ If Q_{q-1} takes on a large value—for example, if it exceeds the 95 percent critical value of the χ_{q-1}^2 distribution—this signals significant serial correlation in returns and suggests that the robust Sharpe ratio, $\widehat{SR}(q)$, should be used instead of $\sqrt{q} \widehat{SR}$ for estimating the Sharpe ratio of q -period returns.

Table 3. Asymptotic Relative Efficiency of Robust Sharpe Ratio Estimator When Returns Are IID

SR	Aggregation Value, q									
	2	3	4	6	12	24	36	48	125	250
0.50	1.06	1.12	1.19	1.34	1.78	2.67	3.56	4.45	10.15	19.41
0.75	1.11	1.24	1.38	1.67	2.54	4.30	6.05	7.81	19.07	37.37
1.00	1.17	1.37	1.58	2.02	3.34	6.00	8.67	11.34	28.45	56.22
1.25	1.22	1.49	1.77	2.34	4.08	7.59	11.09	14.60	37.11	73.66
1.50	1.26	1.59	1.93	2.62	4.72	8.95	13.18	17.42	44.59	88.71
1.75	1.30	1.67	2.06	2.85	5.25	10.08	14.92	19.76	50.81	101.22
2.00	1.33	1.74	2.17	3.04	5.69	11.01	16.34	21.67	55.89	111.45
2.25	1.36	1.80	2.25	3.19	6.04	11.76	17.49	23.23	60.02	119.75
2.50	1.38	1.84	2.33	3.31	6.32	12.37	18.43	24.49	63.38	126.51
2.75	1.40	1.88	2.38	3.42	6.56	12.87	19.20	25.52	66.12	132.02
3.00	1.41	1.91	2.43	3.50	6.75	13.28	19.83	26.37	68.37	136.55

Note: Asymptotic relative efficiency is given by $V_{GMM}(q)/V_{IID}(q)$.

An Empirical Example

To illustrate the potential impact of estimation error and serial correlation in computing Sharpe ratios, I apply the estimators described in the preceding sections to the monthly historical total returns of the 10 largest (as of February 11, 2001) mutual funds from various start dates through June 2000 and 12 hedge funds from various inception dates through December 2000. Monthly total returns for the mutual funds were obtained from the University of Chicago's Center for Research in Security Prices. The 12 hedge funds were selected from the Altvest database to yield a diverse range of annual Sharpe ratios (from 1.00 to 5.00) computed in the standard way ($\sqrt{q}\widehat{SR}$, where \widehat{SR} is the Sharpe ratio estimator applied to monthly returns), with the additional requirement that the funds have a minimum five-year history of returns. The names of the hedge funds have been omitted to maintain their privacy, and I will refer to them only by their investment styles (e.g., relative value fund, risk arbitrage fund).¹¹

Table 4 shows that the 10 mutual funds have little serial correlation in returns, with p -values of Q -statistics ranging from 13.2 percent to 80.2 percent.¹² Indeed, the largest absolute level of autocorrelation among the 10 mutual funds is the 12.4 percent first-order autocorrelation of the Fidelity Magellan Fund. With a risk-free rate of 5/12 percent per month, the monthly Sharpe ratios of the 10 mutual funds range from 0.14 (Growth Fund of America) to 0.32 (Janus Worldwide), with robust standard errors of 0.05 and 0.11, respectively. Because of the lack of serial correlation in the monthly returns of these mutual funds, there is little difference between the IID estimator for the

annual Sharpe ratio, $\sqrt{q}\widehat{SR}$ (in Table 4, $\sqrt{12}\widehat{SR}$), and the robust estimator that accounts for serial correlation, $\widehat{SR}(12)$. For example, even in the case of the Fidelity Magellan Fund, which has the highest first-order autocorrelation among the 10 mutual funds, the difference between a $\sqrt{q}\widehat{SR}$ of 0.73 and a $\widehat{SR}(12)$ of 0.66 is not substantial (and certainly not statistically significant). Note that the robust estimator is marginally lower than the IID estimator, indicating the presence of positive serial correlation in the monthly returns of the Magellan Fund. In contrast, for Washington Mutual Investors, the IID estimate of the annual Sharpe ratio is $\sqrt{q}\widehat{SR} = 0.60$ but the robust estimate is larger, $\widehat{SR}(12) = 0.65$, because of negative serial correlation in the fund's monthly returns (recall that negative serial correlation implies that the variance of the sum of 12 monthly returns is less than 12 times the variance of monthly returns).

The robust standard errors $SE_3(12)$ with $m = 3$ for $\widehat{SR}(12)$ for the mutual funds range from 0.17 (Janus) to 0.47 (Fidelity Growth and Income) and take on similar values when $m = 6$, which indicates that the robust estimator is reasonably well behaved for this dataset. The magnitudes of the standard errors yield 95 percent confidence intervals for annual Sharpe ratios that do not contain 0 for any of the 10 mutual funds. For example, the 95 percent confidence interval for the Vanguard 500 Index fund is $0.85 \pm (1.96 \times 0.26)$, which is (0.33, 1.36). These results indicate Sharpe ratios for the 10 mutual funds that are statistically different from 0 at the 95 percent confidence level.

The results for the 12 hedge funds are different in several respects. The mean returns are higher and the standard deviations lower, implying much

Table 4. Monthly and Annual Sharpe Ratio Estimates for a Sample of Mutual Funds and Hedge Funds

Fund	Start Date	T	$\hat{\mu}$ (%)	$\hat{\sigma}$ (%)	$\hat{\rho}_1$ (%)	$\hat{\rho}_2$ (%)	$\hat{\rho}_3$ (%)	p-Value of Q_{11} (%)	Monthly		Annual			
									\widehat{SR}	SE_3	$\sqrt{12}\widehat{SR}$	$\widehat{SR}(12)$	$SE_3(12)$	$SE_6(12)$
<i>Mutual funds</i>														
Vanguard 500 Index	10/76	286	1.30	4.27	-4.0	-6.6	-4.9	64.5	0.21	0.06	0.72	0.85	0.26	0.25
Fidelity Magellan	1/67	402	1.73	6.23	12.4	-2.3	-0.4	28.6	0.21	0.06	0.73	0.66	0.20	0.21
Investment Company of America	1/63	450	1.17	4.01	1.8	-3.2	-4.5	80.2	0.19	0.05	0.65	0.71	0.22	0.22
Janus	3/70	364	1.52	4.75	10.5	-0.0	-3.7	58.1	0.23	0.06	0.81	0.80	0.17	0.17
Fidelity Contrafund	5/67	397	1.29	4.97	7.4	-2.5	-6.8	58.2	0.18	0.05	0.61	0.67	0.23	0.23
Washington Mutual Investors	1/63	450	1.13	4.09	-0.1	-7.2	-2.6	22.8	0.17	0.05	0.60	0.65	0.20	0.20
Janus Worldwide	1/92	102	1.81	4.36	11.4	3.4	-3.8	13.2	0.32	0.11	1.12	1.29	0.46	0.37
Fidelity Growth and Income	1/86	174	1.54	4.13	5.1	-1.6	-8.2	60.9	0.27	0.09	0.95	1.18	0.47	0.40
American Century Ultra	12/81	223	1.72	7.11	2.3	3.4	1.4	54.5	0.18	0.07	0.64	0.71	0.27	0.25
Growth Fund of America	7/64	431	1.18	5.35	8.5	-2.7	-4.1	45.4	0.14	0.05	0.50	0.49	0.19	0.20
<i>Hedge funds</i>														
Convertible/option arbitrage	5/92	104	1.63	0.97	42.6	29.0	21.4	0.0	1.26	0.28	4.35	2.99	1.04	1.11
Relative value	12/92	97	0.66	0.21	25.9	19.2	-2.1	4.5	1.17	0.17	4.06	3.38	1.16	1.07
Mortgage-backed securities	1/93	96	1.33	0.79	42.0	22.1	16.7	0.1	1.16	0.24	4.03	2.44	0.53	0.54
High-yield debt	6/94	79	1.30	0.87	33.7	21.8	13.1	5.2	1.02	0.27	3.54	2.25	0.74	0.72
Risk arbitrage A	7/93	90	1.06	0.69	-4.9	-10.8	6.9	30.6	0.94	0.20	3.25	3.83	0.87	0.85
Long-short equities	7/89	138	1.18	0.83	-20.2	24.6	8.7	0.1	0.92	0.06	3.19	2.32	0.35	0.37
Multistrategy A	1/95	72	1.08	0.75	48.9	23.4	3.3	0.3	0.89	0.40	3.09	2.18	1.14	1.19
Risk arbitrage B	11/94	74	0.90	0.77	-4.9	2.5	-8.3	96.1	0.63	0.14	2.17	2.47	0.79	0.77
Convertible arbitrage A	9/92	100	1.38	1.60	33.8	30.8	7.9	0.8	0.60	0.18	2.08	1.43	0.44	0.45
Convertible arbitrage B	7/94	78	0.78	0.62	32.4	9.7	-4.5	23.4	0.60	0.18	2.06	1.67	0.68	0.62
Multistrategy B	6/89	139	1.34	1.63	49.0	24.6	10.6	0.0	0.57	0.16	1.96	1.17	0.25	0.25
Fund of funds	10/94	75	1.68	2.29	29.7	21.1	0.9	23.4	0.56	0.19	1.93	1.39	0.67	0.70

Note: For the mutual fund sample, monthly total returns from various start dates through June 2000; for the hedge fund sample, various start dates through December 2000. The term $\hat{\rho}_k$ denotes the k th autocorrelation coefficient, and Q_{11} denotes the Ljung-Box Q -statistic, which is asymptotically χ^2_{11} under the null hypothesis of no serial correlation. \widehat{SR} denotes the usual Sharpe ratio estimator, $(\hat{\mu} - R_f) / \hat{\sigma}$, which is based on monthly data; R_f is assumed to be 5/12 percent per month; and $\widehat{SR}(12)$ denotes the annual Sharpe ratio estimator that takes into account serial correlation in monthly returns. All standard errors are based on GMM estimators using the Newey-West (1982) procedure with truncation lag $m = 3$ for entries in the SE_3 and $SE_3(12)$ columns and $m = 6$ for entries in the $SE_6(12)$ column.

higher Sharpe ratio estimates for hedge funds than for mutual funds. The monthly Sharpe ratio estimates, \widehat{SR} , range from 0.56 (“Fund of funds”) to 1.26 (“Convertible/option arbitrage”), in contrast to the range of 0.14 to 0.32 for the 10 mutual funds. However, the serial correlation in hedge fund returns is also much higher. For example, the first-order autocorrelation coefficient ranges from -20.2 percent to 49.0 percent among the 12 hedge funds, whereas the highest first-order autocorrelation is 12.4 percent among the 10 mutual funds. The p -values provide a more complete summary of the presence of serial

correlation: All but 5 of the 12 hedge funds have p -values less than 5 percent, and several are less than 1 percent.

The impact of serial correlation on the annual Sharpe ratios of hedge funds is dramatic. When the IID estimator, $\sqrt{12}\widehat{SR}$, is used for the annual Sharpe ratio, the “Convertible/option arbitrage” fund has a Sharpe ratio estimate of 4.35, but when serial correlation is properly taken into account by $\widehat{SR}(12)$, the estimate drops to 2.99, implying that the IID estimator overstates the annual Sharpe ratio by 45 percent. The annual Sharpe ratio estimate for the

“Mortgage-backed securities” fund drops from 4.03 to 2.44 when serial correlation is taken into account, implying an overstatement of 65 percent. However, the annual Sharpe ratio estimate for the “Risk arbitrage A” fund *increases* from 3.25 to 3.83 because of negative serial correlation in its monthly returns.

The sharp differences between the annual IID and robust Sharpe ratio estimates underscore the importance of correctly accounting for serial correlation in analyzing the performance of hedge funds. Naively estimating the annual Sharpe ratios by multiplying \widehat{SR} by $\sqrt{12}$ will yield the rank ordering given in the $\widehat{SR} \cdot \sqrt{12}$ column of Table 4, but once serial correlation is taken into account, the rank ordering changes to 3, 2, 5, 7, 1, 6, 8, 4, 10, 9, 12, and 11.

The robust standard errors for the annual robust Sharpe ratio estimates of the 12 hedge funds range from 0.25 to 1.16, which although larger than those in the mutual fund sample, nevertheless imply 95 percent confidence intervals that generally do not include 0. For example, even in the case of the “Multistrategy B” fund, which has the lowest robust Sharpe ratio estimate (1.17), its 95 percent confidence interval is $1.17 \pm 1.96 \times 0.25$, which is (0.68, 1.66). These statistically significant Sharpe ratios are consistent with previous studies that document the fact that hedge funds do seem to exhibit statistically significant excess returns.¹³ The similarity of the standard errors between the $m = 3$ and $m = 6$ cases for the hedge fund sample indicates that the robust estimator is also well behaved in this case, despite the presence of significant serial correlation in monthly returns.

In summary, the empirical examples illustrate the potential impact that serial correlation can have on Sharpe ratio estimates and the importance of properly accounting for departures from the standard IID framework. Robust Sharpe ratio estimators contain significant additional information about the risk–reward trade-offs for active investment products, such as hedge funds; more detailed analysis of the risks and rewards of hedge fund investments is performed in Getmansky, Lo, and Makarov (2002) and Lo (2001).

Conclusion

Although the Sharpe ratio has become part of the canon of modern financial analysis, its applications typically do not account for the fact that it is an estimated quantity, subject to estimation errors that can be substantial in some cases. The results presented in this article provide one way to gauge the accuracy of these estimators, and it should come as no surprise that the statistical properties of Sharpe ratios depend intimately on the statistical properties of the return series on which they are based. This suggests that a more sophisticated approach to interpreting Sharpe ratios is called for, one that incorporates information about the investment style that generates the returns and the market environment in which those returns are generated. For example, hedge funds have very different return characteristics from the characteristics of mutual funds; hence, the comparison of Sharpe ratios between these two investment vehicles cannot be performed naively. In light of the recent interest in alternative investments by institutional investors—investors that are accustomed to standardized performance attribution measures such as the annualized Sharpe ratio—there is an even greater need to develop statistics that are consistent with a portfolio’s investment style.

The empirical example underscores the practical relevance of proper statistical inference for Sharpe ratio estimators. Ignoring the impact of serial correlation in hedge fund returns can yield annualized Sharpe ratios that are overstated by more than 65 percent, understated Sharpe ratios in the case of negatively serially correlated returns, and inconsistent rankings across hedge funds of different styles and objectives. By using the appropriate statistical distribution for quantifying the performance of each return history, the Sharpe ratio can provide a more complete understanding of the risks and rewards of a broad array of investment opportunities.

I thank Nicholas Chan, Arnout Eikeboom, Jim Holub, Chris Jakob, Laurel Kenner, Frank Linet, Jon Markman, Victor Niederhoffer, Dan O’Reilly, Bill Sharpe, and Jonathan Taylor for helpful comments and discussion. Research support from AlphaSimplex Group is gratefully acknowledged.

Appendix A. Asymptotic Distributions of Sharpe Ratio Estimators

The first section of this appendix presents results for IID returns, and the second section presents corresponding results for non-IID returns. Results for time-aggregated Sharpe ratios are reported in the third section, and in the final section, the asymptotic variance, $V_{GMM}(q)$, of the time-aggregated robust estimator, $\widehat{SR}(q)$, is derived for the special case of IID returns.

Throughout the appendix, the following conventions are maintained: (1) all vectors are column vectors unless otherwise indicated; (2) vectors and matrixes are always typeset in boldface (i.e., X and μ are scalars and \mathbf{X} and $\boldsymbol{\mu}$ are vectors or matrixes).

IID Returns

To derive an expression for the asymptotic distribution of \widehat{SR} , we must first obtain the asymptotic *joint* distribution of $\hat{\mu}$ and $\hat{\sigma}^2$. Denote by $\hat{\boldsymbol{\theta}}$ the column vector $(\hat{\mu} \ \hat{\sigma}^2)'$ and let $\boldsymbol{\theta}$ denote the corresponding column vector of population values $(\mu \ \sigma^2)'$. If returns are IID, it is a well-known consequence of the Central Limit Theorem that the asymptotic distribution of $\hat{\boldsymbol{\theta}}$ is given by (see White):

$$\sqrt{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \stackrel{d}{\approx} N(0, \mathbf{V}_{\boldsymbol{\theta}}), \mathbf{V}_{\boldsymbol{\theta}} \equiv \begin{pmatrix} \sigma^2 & 0 \\ 0 & 2\sigma^4 \end{pmatrix}, \quad (\text{A1})$$

where the notation $\stackrel{d}{\approx}$ indicates that this is an asymptotic approximation. Because the Sharpe ratio estimator \widehat{SR} can be written as a function $g(\hat{\boldsymbol{\theta}})$ of $\hat{\boldsymbol{\theta}}$, its asymptotic distribution follows directly from Taylor's theorem or the so-called delta method (see, for example, White):

$$\sqrt{T}[g(\hat{\boldsymbol{\theta}}) - g(\boldsymbol{\theta})] \stackrel{d}{\approx} N(0, \mathbf{V}_g), \mathbf{V}_g \equiv \frac{\partial g}{\partial \boldsymbol{\theta}} \mathbf{V}_{\boldsymbol{\theta}} \frac{\partial g}{\partial \boldsymbol{\theta}'}. \quad (\text{A2})$$

In the case of the Sharpe ratio, $g(\cdot)$ is given by Equation 2; hence,

$$\frac{\partial g}{\partial \boldsymbol{\theta}'} = \begin{bmatrix} 1/\sigma \\ -(\mu - R_f)/(2\sigma^3) \end{bmatrix}, \quad (\text{A3})$$

which yields the following asymptotic distribution for \widehat{SR} :

$$\sqrt{T}(\widehat{SR} - SR) \stackrel{d}{\approx} N(0, V_{IID}), V_{IID} = 1 + \frac{(\mu - R_f)^2}{2\sigma^2} = 1 + \frac{1}{2} SR^2. \quad (\text{A4})$$

Non-IID Returns

Denote by \mathbf{X}_t the vector of period- t returns and lags $(R_t R_{t-1} \dots R_{t-q+1})'$ and let (\mathbf{X}_t) be a stochastic process that satisfies the following conditions:

- H1: $\{\mathbf{X}_t : t \in (-\infty, \infty)\}$ is stationary and ergodic;
- H2: $\boldsymbol{\theta}_0 \in \Theta$, Θ is an open subset of \mathfrak{R}^k ;
- H3: $\forall \boldsymbol{\theta} \in \Theta$, $\varphi(\cdot, \boldsymbol{\theta})$ and $\varphi_{\boldsymbol{\theta}}(\cdot, \boldsymbol{\theta})$ are Borel measurable and $\varphi_{\boldsymbol{\theta}}(\mathbf{X}, \cdot)$ is continuous on Θ for all \mathbf{X} ;
- H4: $\varphi_{\boldsymbol{\theta}}$ is first-moment continuous at $\boldsymbol{\theta}_0$; $E[\varphi_{\boldsymbol{\theta}}(\mathbf{X}, \cdot)]$ exists, is finite, and is of full rank.
- H5: Let $\varphi_t \equiv \varphi(\mathbf{X}_t, \boldsymbol{\theta}_0)$

and

$$\mathbf{v}_j \equiv E[\varphi_0 | \varphi_{-1}, \varphi_{-2}, \dots] - E[\varphi_0 | \varphi_{-j-1}, \varphi_{-j-2}, \dots]$$

and assume

- (i): $E[\boldsymbol{\varphi}_0 \boldsymbol{\varphi}'_0]$ exists and is finite,
- (ii): \mathbf{v}_j converges in mean square to 0, and
- (iii): $\sum_{j=0}^{\infty} E(\mathbf{v}'_j \mathbf{v}_j)^{1/2}$ is finite,

which implies $E[\boldsymbol{\varphi}(\mathbf{X}_t, \boldsymbol{\theta}_0)] = 0$.

$$\text{H6: Let } \hat{\boldsymbol{\theta}} \text{ solve } \frac{1}{T} \sum_{t=1}^T \boldsymbol{\varphi}(\mathbf{X}_t, \boldsymbol{\theta}) = 0.$$

Then, Hansen shows that

$$\sqrt{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \stackrel{d}{\sim} N(0, \mathbf{V}_{\boldsymbol{\theta}}), \mathbf{V}_{\boldsymbol{\theta}} \equiv \mathbf{H}^{-1} \boldsymbol{\Sigma} \mathbf{H}^{-1'} \quad (\text{A5})$$

where

$$\mathbf{H} \equiv \lim_{T \rightarrow \infty} E \left[\frac{1}{T} \sum_{t=1}^T \boldsymbol{\varphi}_{\boldsymbol{\theta}}(\mathbf{X}_t, \boldsymbol{\theta}_0) \right], \quad (\text{A6})$$

$$\boldsymbol{\Sigma} \equiv \lim_{T \rightarrow \infty} E \left[\frac{1}{T} \sum_{t=1}^T \sum_{s=1}^T \boldsymbol{\varphi}(\mathbf{X}_t, \boldsymbol{\theta}_0) \boldsymbol{\varphi}(\mathbf{X}_s, \boldsymbol{\theta}_0)' \right], \quad (\text{A7})$$

and $\boldsymbol{\varphi}_{\boldsymbol{\theta}}(R_t, \boldsymbol{\theta})$ denotes the derivative of $\boldsymbol{\varphi}(R_t, \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$.¹⁴ Specifically, let $\boldsymbol{\varphi}(R_t, \boldsymbol{\theta})$ denote the following vector function:

$$\boldsymbol{\varphi}(R_t, \boldsymbol{\theta}) \equiv \begin{bmatrix} R_t - \mu \\ (R_t - \mu)^2 - \sigma^2 \end{bmatrix}. \quad (\text{A8})$$

The GMM estimator of $\boldsymbol{\theta}$, denoted by $\hat{\boldsymbol{\theta}}$, is given implicitly by the solution to

$$\frac{1}{T} \sum_{t=1}^T \boldsymbol{\varphi}(R_t, \boldsymbol{\theta}) = 0, \quad (\text{A9})$$

which yields the standard estimators $\hat{\mu}$ and $\hat{\sigma}^2$ given in Equation 3. For the moment conditions in Equation A8, \mathbf{H} is given by:

$$\mathbf{H} \equiv \lim_{T \rightarrow \infty} E \left\{ \frac{1}{T} \sum_{t=1}^T \begin{bmatrix} -1 & 0 \\ 2(\mu - R_t) & -1 \end{bmatrix} \right\} = -\mathbf{I}. \quad (\text{A10})$$

Therefore, $\mathbf{V}_{\boldsymbol{\theta}} = \boldsymbol{\Sigma}$ and the asymptotic distribution of the Sharpe ratio estimator follows from the delta method as in the first section:

$$\sqrt{T}(\widehat{\text{SR}} - \text{SR}) \stackrel{d}{\sim} N(0, V_{GMM}), V_{GMM} = \frac{\partial g}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma} \frac{\partial g}{\partial \boldsymbol{\theta}'}, \quad (\text{A11})$$

where $\partial g / \partial \boldsymbol{\theta}$ is given in Equation A3. An estimator for $\partial g / \partial \boldsymbol{\theta}$ may be obtained by substituting $\hat{\boldsymbol{\theta}}$ into Equation A3, and an estimator for $\boldsymbol{\Sigma}$ may be obtained by using Newey and West's (1987) procedure:

$$\hat{\boldsymbol{\Sigma}} = \hat{\boldsymbol{\Omega}}_0 + \sum_{j=1}^m \omega(j, m) (\hat{\boldsymbol{\Omega}}_j + \hat{\boldsymbol{\Omega}}'_j), \quad m \ll T, \quad (\text{A12})$$

$$\hat{\boldsymbol{\Omega}}_j \equiv \frac{1}{T} \sum_{t=j+1}^T \boldsymbol{\varphi}(R_t, \hat{\boldsymbol{\theta}}) \boldsymbol{\varphi}(R_{t-j}, \hat{\boldsymbol{\theta}})', \quad (\text{A13})$$

$$\omega(j, m) \equiv 1 - \frac{j}{m+1}, \quad (\text{A14})$$

and m is the truncation lag, which must satisfy the condition $m/T \rightarrow \infty$ as T increases without bound to ensure consistency. An estimator for V_{SR} can then be constructed as

$$\hat{V}_{GMM} = \frac{\partial g(\hat{\theta})}{\partial \theta} \hat{\Sigma} \frac{\partial g(\hat{\theta})}{\partial \theta}'. \tag{A15}$$

Time Aggregation

Let $\theta \equiv [\mu \ \sigma^2 \ \gamma_1 \ \dots \ \gamma_{q-1}]'$ denote the vector of parameters to be estimated, where γ_k is the k th-order autocovariance of R_t , and define the following moment conditions:

$$\begin{aligned} \varphi_1(\mathbf{X}_t, \theta) &= R_t - \mu \\ \varphi_2(\mathbf{X}_t, \theta) &= (R_t - \mu)^2 - \sigma^2 \\ \varphi_3(\mathbf{X}_t, \theta) &= (R_t - \mu)(R_{t-1} - \mu) - \gamma_1 \\ \varphi_4(\mathbf{X}_t, \theta) &= (R_t - \mu)(R_{t-2} - \mu) - \gamma_2 \\ &\vdots \\ \varphi_{q+1}(\mathbf{X}_t, \theta) &= (R_t - \mu)(R_{t-q+1} - \mu) - \gamma_{q-1} \\ \boldsymbol{\varphi}(\mathbf{X}_t, \theta) &\equiv [\varphi_1 \ \varphi_2 \ \varphi_3 \ \dots \ \varphi_{q+1}]', \end{aligned} \tag{A16}$$

where $\mathbf{X}_t \equiv [R_t \ R_{t-1} \ \dots \ R_{t-q+1}]'$. The GMM estimator $\hat{\theta}$ is defined by Equation A9, which yields the standard estimators $\hat{\mu}$ and $\hat{\sigma}^2$ in Equation 3 as well as the standard estimators for the autocovariances:

$$\hat{\gamma}_k = \frac{1}{T} \sum_{t=k+1}^T (R_t - \hat{\mu})(R_{t-k} - \hat{\mu}). \tag{A17}$$

The estimator for the Sharpe ratio then follows directly:

$$\widehat{SR}(q) = \hat{\eta}(q) \widehat{SR}, \quad \hat{\eta}(q) \equiv \frac{q}{\sqrt{q + 2 \sum_{k=1}^{q-1} (q-k) \hat{\rho}_k}}, \tag{A18}$$

where

$$\hat{\rho}_k = \frac{\hat{\gamma}_k}{\hat{\sigma}^2}.$$

As in the first two sections of this appendix, the asymptotic distribution of $\widehat{SR}(q)$ can be obtained by applying the delta method to $g(\hat{\theta})$ where the function $g(\cdot)$ is now given by Equation 20. Recall from Equation A5 that the asymptotic distribution of the GMM estimator $\hat{\theta}$ is given by

$$\sqrt{T}(\hat{\theta} - \theta) \stackrel{d}{\sim} N(0, \mathbf{V}_\theta), \quad \mathbf{V}_\theta \equiv \mathbf{H}^{-1} \boldsymbol{\Sigma} \mathbf{H}^{-1}, \tag{A19}$$

$$\boldsymbol{\Sigma} \equiv \lim_{T \rightarrow \infty} \mathbf{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{s=1}^T \boldsymbol{\varphi}(\mathbf{X}_t, \theta) \boldsymbol{\varphi}(\mathbf{X}_s, \theta)' \right]. \tag{A20}$$

For the moment conditions in Equation A16, \mathbf{H} is

$$\mathbf{H} = \lim_{T \rightarrow \infty} \mathbf{E} \left\{ \frac{1}{T} \sum_{t=1}^T \begin{bmatrix} -1 & 0 & 0 & \dots & 0 \\ 2(\mu - R_t) & -1 & 0 & \dots & 0 \\ 2\mu - R_t - R_{t-1} & 0 & -1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 2\mu - R_t - R_{t-q+1} & 0 & \dots & 0 & -1 \end{bmatrix} \right\} = -\mathbf{I}; \tag{A21}$$

hence, $\mathbf{V}_\theta = \Sigma$. The asymptotic distribution of $\widehat{\text{SR}}(q)$ then follows from the delta method:

$$\sqrt{T}[\widehat{\text{SR}}(q) - \text{SR}(q)] \stackrel{d}{\sim} N\left[0, V_{\text{GMM}}(q)\right], V_{\text{GMM}}(q) = \frac{\partial g}{\partial \theta} \Sigma \frac{\partial g}{\partial \theta}', \quad (\text{A22})$$

where the components of $\partial g / \partial \theta$ are

$$\frac{\partial g}{\partial \mu} = \frac{q}{\sigma \sqrt{q + 2 \sum_{k=1}^{q-1} (q-k) \rho_k}} \quad (\text{A23})$$

$$\frac{\partial g}{\partial \sigma^2} = -\frac{q^2 \text{SR}}{2\sigma^2 \left[q + 2 \sum_{k=2}^{q-1} (q-k) \rho_k \right]^{3/2}}, \quad (\text{A24})$$

$$\frac{\partial g}{\partial \gamma_k} = -\frac{q(q-k) \text{SR}}{\sigma^2 \left[q + 2 \sum_{k=1}^{q-1} (q-k) \rho_k \right]^{3/2}}, \quad k=1, \dots, q-1, \quad (\text{A25})$$

and

$$\frac{\partial g}{\partial \theta} = \left[\frac{\partial g}{\partial \mu} \quad \frac{\partial g}{\partial \sigma^2} \quad \frac{\partial g}{\partial \gamma_1} \quad \dots \quad \frac{\partial g}{\partial \gamma_{q-1}} \right]. \quad (\text{A26})$$

Substituting $\hat{\theta}$ into Equation A26, estimating Σ according to Equation A12, and forming the matrix product $\partial g(\hat{\theta}) / \partial \theta \widehat{\Sigma} \partial g(\hat{\theta}) / \partial \theta'$ yields an estimator for the asymptotic variance of $\widehat{\text{SR}}(q)$.

Using $\widehat{V}_{\text{GMM}}(q)$ When Returns Are IID

For IID returns, it is possible to evaluate Σ in Equation A20 explicitly as

$$\Sigma = \begin{pmatrix} \sigma^2 & v_3 & 0 & 0 & \dots & 0 \\ v_3 & v_4 - \sigma^4 & 0 & 0 & \dots & 0 \\ 0 & 0 & \sigma^4 & 0 & \dots & \vdots \\ 0 & 0 & 0 & \sigma^4 & \dots & 0 \\ 0 & 0 & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \sigma^4 \end{pmatrix} = \begin{pmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{pmatrix}, \quad (\text{A27})$$

where $v_3 \equiv E[(R_t - \mu)^3]$, $v_4 \equiv E[(R_t - \mu)^4]$, and Σ is partitioned into a block-diagonal matrix with a (2×2) matrix Σ_1 and a diagonal $(q-1) \times (q-1)$ matrix $\Sigma_2 = \sigma^4 \mathbf{I}$ along its diagonal. Because $\gamma_k = 0$ for all $k > 0$, $\partial g / \partial \theta$ simplifies to

$$\frac{\partial g}{\partial \mu} = \frac{\sqrt{q}}{\sigma}, \quad (\text{A28})$$

$$\frac{\partial g}{\partial \sigma^2} = -\frac{\sqrt{q} \text{SR}}{2\sigma^2}, \quad (\text{A29})$$

$$\frac{\partial g}{\partial \gamma_k} = -\frac{q(q-k) \text{SR}}{\sigma^2 q^{3/2}}, \quad k=1, \dots, q-1, \quad (\text{A30})$$

and

$$\frac{\partial g}{\partial \theta} = \left(\frac{\partial g}{\partial \mu} \quad \frac{\partial g}{\partial \sigma^2} \quad \frac{\partial g}{\partial \gamma_1} \quad \dots \quad \frac{\partial g}{\partial \gamma_{q-1}} \right) = [\mathbf{a} \quad \mathbf{b}], \quad (\text{A31})$$

where $\partial g/\partial \theta$ is also partitioned to conform to the partitioned matrix Σ in Equation A27. Therefore, the asymptotic variance of the robust estimator $\widehat{SR}(q)$ is given by

$$\begin{aligned}
 V_{GMM}(q) &= \frac{\partial g}{\partial \theta} \Sigma \frac{\partial g}{\partial \theta}' = \mathbf{a} \Sigma_1 \mathbf{a}' + \mathbf{b} \Sigma_2 \mathbf{b}' \\
 &= q \left[1 - \frac{v_3 \text{SR}}{\sigma^3} + (v_4 - \sigma^4) \frac{\text{SR}^2}{4\sigma^4} \right] + (\sqrt{q} \text{SR})^2 \sum_{j=1}^{q-1} \left(1 - \frac{j}{q} \right)^2.
 \end{aligned} \tag{A32}$$

If R_t is normally distributed, then $v_3 = 0$ and $v_4 = 3\sigma^4$; hence,

$$\begin{aligned}
 V_{GMM}(q) &= q \left[1 + (3\sigma^4 - \sigma^4) \frac{\text{SR}^2}{4\sigma^4} \right] + (\sqrt{q} \text{SR})^2 \sum_{j=1}^{q-1} \left(1 - \frac{j}{q} \right)^2 \\
 &= q \left(1 + \frac{1}{2} \text{SR}^2 \right) + (\sqrt{q} \text{SR})^2 \sum_{j=1}^{q-1} \left(1 - \frac{j}{q} \right)^2 \\
 &= V_{IID}(q) + (\sqrt{q} \text{SR})^2 \sum_{j=1}^{q-1} \left(1 - \frac{j}{q} \right)^2 \geq V_{IID}(q).
 \end{aligned} \tag{A33}$$

Notes

1. See Sharpe (1994) for an excellent review of its many applications, as well as some new extensions.
2. The Central Limit Theorem is a remarkable mathematical discovery on which much of modern statistical inference is based. It shows that under certain conditions, the probability distribution of a properly normalized sum of random variables must converge to the standard normal distribution, regardless of how each of the random variables in the sum is distributed. Therefore, using the normal distribution for calculating significance levels and confidence intervals is often an excellent approximation, even if normality does not hold for the particular random variables in question. See White (1984) for a rigorous exposition of the role of the Central Limit Theorem in modern econometrics.
3. See, for example, Lo and MacKinlay (1999) and their citations.
4. Additional regularity conditions are required; see Appendix A, Hansen (1982), and White for further discussion.
5. The term “robust” is meant to convey the ability of an estimator to perform well under various sets of assumptions. Another commonly used term for such estimators is “nonparametric,” which indicates that an estimator is not based on any parametric assumption, such as normally distributed returns. See Randles and Wolfe (1979) for further discussion of nonparametric estimators and Hansen for the generalized method of moments estimator.
6. See, for example, Campbell, Lo, and MacKinlay (1997, Ch. 9), Lo and MacKinlay (Ch. 4), Merton (1980), and Shiller and Perron (1985).
7. The exact expression is, of course,

$$R_t(q) \equiv \prod_{j=0}^{q-1} (1 + R_{t-j}) - 1.$$
 For most (but not all) applications, Equation 16 is an excellent approximation. Alternatively, if R_t is defined to be the continuously compounded return [i.e., $R_t \equiv \log(P_t/P_{t-1})$, where P_t is the price or net asset value at time t], then Equation 16 is exact.
8. See Bodie (1995) and the ensuing debate regarding risks in the long run for further evidence of the inadequacy of the Sharpe ratio—or any other single statistic—for delineating the risk–reward profile of a dynamic investment policy.
9. The k th-order autocorrelation of a time series R_t is defined as the correlation coefficient between R_t and R_{t-k} , which is simply the covariance between R_t and R_{t-k} divided by the square root of the product of the variances of R_t and R_{t-k} . But because the variances of R_t and R_{t-k} are the same under our assumption of stationarity, the denominator of the autocorrelation is simply the variance of R_t .
10. See, for example, Harvey (1981, Ch. 6.2).
11. These are the investment styles reported in the Altvest database; no attempt was made to verify or to classify the hedge funds independently.
12. The p -value of a statistic is defined as the smallest level of significance for which the null hypothesis can be rejected based on the statistic’s value. In particular, the p -value of 16.0 percent for the Q -statistic of Washington Mutual Investors in Table 4 implies that the null hypothesis of no serial correlation can be rejected only at the 16.0 percent significance level; at any lower level of significance—say, 5 percent—the null hypothesis cannot be rejected. Therefore, smaller p -values indicate stronger evidence against the null hypothesis and larger p -values indicate stronger evidence in favor of the null. Researchers often report p -values instead of test statistics because p -values are easier to interpret. To interpret a test statistic, one must compare it with the critical values of the appropriate distribution. This comparison is performed in computing the p -value. For further discussion of p -values and their interpretation, see, for example, Bickel and Doksum (1977, Ch. 5.2.B).
13. See, for example, Ackermann, McEnally, and Ravenscraft (1999), Brown, Goetzmann, and Ibbotson (1999), Brown, Goetzmann, and Park (2001), Fung and Hsieh (1997a, 1997b, 2000), and Liang (1999, 2000, 2001).
14. See Magnus and Neudecker (1988) for the specific definitions and conventions of vector and matrix derivatives of vector functions.

References

- Ackermann, C., R. McEnally, and D. Ravenscraft. 1999. “The Performance of Hedge Funds: Risk, Return, and Incentives.” *Journal of Finance*, vol. 54, no. 3 (June):833–874.
- Bickel, P., and K. Doksum. 1977. *Mathematical Statistics: Basic Ideas and Selected Topics*. San Francisco, CA: Holden-Day.
- Bodie, Z. 1995. “On the Risk of Stocks in the Long Run.” *Financial Analysts Journal*, vol. 51, no. 3 (May/June):18–22.
- Brown, S., W. Goetzmann, and R. Ibbotson. 1999. “Offshore Hedge Funds: Survival and Performance 1989–1995.” *Journal of Business*, vol. 72, no. 1 (January):91–118.
- Brown, S., W. Goetzmann, and J. Park. 1997. “Careers and Survival: Competition and Risks in the Hedge Fund and CTA Industry.” *Journal of Finance*, vol. 56, no. 5 (October):1869–86.
- Campbell, J., A. Lo, and C. MacKinlay. 1997. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.
- Fung, W., and D. Hsieh. 1997a. “Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds.” *Review of Financial Studies*, vol. 10, no. 2 (April):75–302.
- . 1997b. “Investment Style and Survivorship Bias in the Returns of CTAs: The Information Content of Track Records.” *Journal of Portfolio Management*, vol. 24, no. 1 (Fall):30–41.
- . 2000. “Performance Characteristics of Hedge Funds and Commodity Funds: Natural versus Spurious Biases.” *Journal of Financial and Quantitative Analysis*, vol. 35, no. 3 (September):291–307.
- Getmansky, M., A. Lo, and I. Makarov. 2002. “An Econometric Model of Illiquidity and Performance Smoothing in Hedge Fund Returns.” Unpublished manuscript, MIT Laboratory for Financial Engineering.
- Hansen, L. 1982. “Large Sample Properties of Generalized Method of Moments Estimators.” *Econometrica*, vol. 50, no. 4 (July):1029–54.
- Harvey, A. 1981. *Time Series Models*. New York: John Wiley & Sons.
- Liang, B. 1999. “On the Performance of Hedge Funds.” *Financial Analysts Journal*, vol. 55, no. 4 (July/August):72–85.

- . 2000. "Hedge Funds: The Living and the Dead." *Journal of Financial and Quantitative Analysis*, vol. 35, no. 3 (September):309–326.
- . 2001. "Hedge Fund Performance: 1990–1999." *Financial Analysts Journal*, vol. 57, no. 1 (January/February):11–18.
- Ljung, G., and G. Box. 1978. "On a Measure of Lack of Fit in Time Series Models." *Biometrika*, vol. 65, no. 1: 297–303.
- Lo, A. 2001. "Risk Management for Hedge Funds: Introduction and Overview." *Financial Analysts Journal*, vol. 57, no. 6 (November/December):16–33.
- Lo, A., and C. MacKinlay. 1999. *A Non-Random Walk Down Wall Street*. Princeton, NJ: Princeton University Press.
- Magnus, J., and H. Neudecker. 1988. *Matrix Differential Calculus: With Applications in Statistics and Economics*. New York: John Wiley & Sons.
- Merton, R. 1980. "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics*, vol. 8, no. 4 (December):323–361.
- Newey, W., and K. West. 1987. "A Simple Positive Definite Heteroscedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, vol. 55, no. 3:703–705.
- Randles, R., and D. Wolfe. 1979. *Introduction to the Theory of Nonparametric Statistics*. New York: John Wiley & Sons.
- Sharpe, W. 1994. "The Sharpe Ratio." *Journal of Portfolio Management*, vol. 21, no. 1 (Fall):49–58.
- Shiller, R., and P. Perron. 1985. "Testing the Random Walk Hypothesis: Power versus Frequency of Observation." *Economics Letters*, vol. 18, no. 4:381–386.
- White, H. 1984. *Asymptotic Theory for Econometricians*. New York: Academic Press.