Financial Analysts Journal ›

Volume 58, 2002 - Issue 4



PERFORMANCE MEASUREMENT AND EVALUATION

# The Statistics of Sharpe Ratios

Andrew W. Lo

Pages 36-52 | Published online: 02 Jan 2019

Cite this article https://doi.org/10.2469/faj.v58.n4.2453



## Abstract

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.

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? In this article, I provide an answer by deriving the statistical distributions of the usual Sharpe ratio estimator—sample mean excess return over sample standard deviation—using standard econometric methods under several different sets of assumptions for the statistical properties of the return series on which the ratio is based. Armed with these statistical distributions, I show that confidence intervals, standard errors, and hypothesis tests for the estimated Sharpe ratio can be computed 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 (e.g., time-series properties, such as mean reversion, momentum, or time-varying volatilities). Although this may seem like a theoretical exercise best left for statisticians, 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).

At a superficial level, it is obvious that the properties of the Sharpe ratio should depend on the investment style of the portfolio being evaluated; 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 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 when computing an annualized Sharpe ratio from monthly data. For example, 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 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.

Although the Sharpe ratio has become part of the canon of modern financial analysis, the results presented in this article suggest that a more sophisticated approach to interpreting Sharpe ratios is called for, one that incorporates information about the investment style that generated the returns and the market environment in which those returns were 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 in this article underscores the practical relevance of proper statistical inference for Sharpe ratio estimators: Ignoring the impact of serial correlation of 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.