BETTER VARIANCE MEASUREMENT

Gauging the margin of error in mean-variance optimization.

MARK KRITZMAN

Managing partner, Windham Capital Management, for State Street Global Markets.

Investors have long worried about the approximation error of mean-variance optimization, which relies on restrictive assumptions about investor utility and return distributions. Levy and Markowitz have previously shown that mean-variance approximations to power utility are remarkably accurate, based on a sample of mutual fund returns. Samuelson proposes a new metric, which he calls gratuitous dead weight loss, to gauge the accuracy of mean-variance optimization in the presence of non-quadratic utility and non-normality. This issue is critically important to many investors given the widespread use of hedge funds, which often display significant departures from normality.

We measured the gratuitous dead weight loss of mean-variance optimization for a variety of plausible utility functions based on representative empirical distributions. For the most part, the gratuitous dead weight loss of mean-variance optimization is negligible—only a fraction of a basis point.

"FOR POWER UTILITY, MEAN-VARIANCE OPTIMIZATION PERFORMS WELL BY VIRTUALLY ANY STANDARD."

We also measured the gratuitous dead weight loss associated with sampling error for both mean-variance optimization and full-sample utility maximization. We found that the gratuitous dead weight loss associated with sampling error for both mean-variance optimization and full-sample utility maximization is still quite small, but more than 10 times as great as the gratuitous dead weight loss associated with mean-variance approximation error.

We next introduced a hybrid approach to portfolio formation in which we simulated theoretically based non-normal distributions. We then used the entire sample to maximize expected utility based on plausible utility functions. Our simulations suggest that, for power utility, mean-variance optimization performs well by virtually any standard.

Finally, we introduced a bilinear utility function and applied mean-variance optimization to a sample of undiversified hedge funds. In this case, mean-variance optimization performed poorly.

We offer three caveats. First, we assumed throughout our analyses that investors have some variation of power utility. Even though mean-variance optimization produces portfolio weights that vary by as much as 20% from full-sample utility maximization in our examples with manufactured skewness and kurtosis, their certainty equivalents are not particularly different. Higher moments seem not to matter very much for investors with power utility. However, some investors might have very different utility functions. Perhaps certain investors are faced with thresholds which, if they were to breach, would sharply reduce their utility. Given this type of bilinear utility function, mean-variance optimization generates substantial gratuitous dead weight loss.

Second, we have only considered allocation among a few assets of moderate risk. Mean-variance optimization may not perform as well when applied to portfolios comprising a large number of relatively risky securities, especially if these securities are highly correlated.

Third, even if we choose to ignore higher moments for the purpose of forming portfolios, we should not ignore them when we assess exposure to loss or manager skill. Value-at-Risk, for example, is typically much higher for portfolios with significant negative skewness and kurtosis than we would infer from a normal distribution. In addition, some investment strategies may appear to add value based on mean and variance, when in fact the apparent value added is an artifact of higher moments. Higher moments do matter, but typically not for choosing portfolios. ∎