NEW FRONTIER

Are Good Estimates Enough? No.

by

Richard O. Michaud and Robert O. Michaud

IMCA
January 2009

The mean-variance (MV) optimization of Markowitz (1959) has been the standard for efficient asset allocation for almost 50 years. Nearly all commercial asset allocation optimizers are based on some variation of the Markowitz method. Markowitz MV optimized portfolios potentially have many attractive investment characteristics. Optimized portfolios may reduce risk without reducing expected return. MV optimization also enables tailoring portfolios to various risk and return preferences.

However, even the best risk and return estimates are uncertain. While theoretically important for modern finance, MV optimization does not adjust for uncertainty. As a result, MV optimization typically results in an unstable process, unintuitive optimized portfolios, and poor out-of-sample performance. Tests demonstrate that equally weighted portfolios often prevail over MV optimized portfolios and that MV optimized portfolios may have little, if any, investment value.[i] These limitations result primarily from the way investment information is used in MV optimization.

Many institutions take for granted the characteristics of their optimization technology. Their implicit assumption is that good risk and return inputs are all that matter in defining an optimal portfolio. Investment institutions focus the bulk of their resources on improving the reliability of their forecasts of risk and return. Academics frequently address the acknowledged limitations of MV optimization in practice by proposing new ways to improve the inputs, such as Bayesian estimation. While these sophisticated proposals may improve the reliability of investment estimates, the enhanced inputs often do not overcome the investment limitations of MV optimization. Reliable inputs are certainly important, but even excellent risk and return inputs do not have the perfect certainty required of MV optimization. MV optimization primarily is used as a convenient framework for imposing ad hoc constraints and as a scientific veneer to the asset allocation process.

The Resampled Efficient Frontier™ (REF) or Resampled Efficiency™ (RE) optimization is a generalization of the MV efficient frontier that leads to superior investment performance on average.[ii] RE optimization addresses forecast uncertainty with Monte Carlo methods. RE optimization is essentially a means for controlling the level of certainty in investment information in the MV optimization process. As we will show, this statistical approach to defining portfolio optimality is provably effective at improving optimized portfolio performance on average.

The Test[iii]

Suppose you have found good estimates of future risk and return. What do you hope to see in your investment performance? You may expect that the efficient portfolio you compute from your information is roughly, on average, what you observe in the investment period. We develop such a set of inputs and then compare how they are used in MV and RE optimization. Figure 1 displays the results of our simulation tests. The solid black curve depicts the true MV efficient frontier.[iv]

FIGURE 1: IN- AND OUT-OF-SAMPLE TESTS

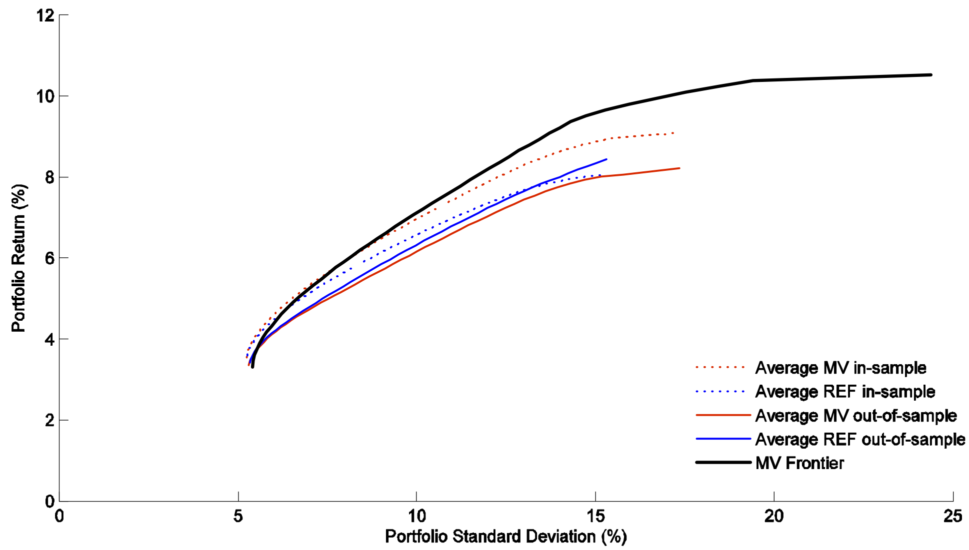

Investors know that there always is uncertainty in investment forecasts. A way to model your uncertainty about the true risks and returns is to use Monte Carlo simulation or resampling to compute new estimates by simulating returns. For example, a Monte Carlo simulation of 100 returns for an asset with an assumed 10% expected return and 20% standard deviation will result in a mean and standard deviation different from the 10% and 20% inputs. These Monte Carlo estimates of risk and return have estimation error relative to the original estimates and help to quantify the effect of the uncertainty inherent in investment information on the optimization process.

From the simulated returns, New Frontier computes new optimization inputs and the associated simulated MV efficient frontier. We use the simulated inputs to compute REF optimal portfolios. We base the REF on an additional set of Monte Carlo simulations of estimated inputs and efficient frontiers and a patented averaging process. REF optimality is defined by averaging the many ways things can happen that are consistent with what you think you know. The resulting REF portfolios are better diversified than their corresponding MV efficient portfolios because their construction considers many more alternative investment scenarios.

In practice, investors typically modify historical average returns and risk estimates in order to enhance forecast value. To mirror this process, we use simple yet powerful statistical procedures, called James-Stein and Ledoit estimation, that are known to improve the forecast value of historically estimated risk-return inputs on average.[v] After the Monte Carlo process simulates returns, we use the new estimates of the optimization inputs for both the MV and RE portfolios. In effect, these are designed to be very "good" inputs from a modern statistical point of view.[vi]

The test proceeds by computing many simulated MV efficient frontiers, each statistically consistent with your original forecasts and level of uncertainty and averaging the results. The average of the simulated MV efficient frontiers is displayed as the red dashed curve in figure 1. The average of all the associated simulated REFs is displayed as the blue dashed curve in figure 1.

The red and blue dashed curves represent averages of in-sample efficient frontiers. Because estimation error always exists in practice, the dashed curves represent the efficient frontiers you see when you invest with either MV or RE optimization. But because this is a simulation, we can go back to the original data (the inputs behind the black curve) and see how estimation error led to misestimation of the original efficient frontier. This is the basis of the out-of-sample test of the optimization process. The solid red curve represents the actual out-of-sample average risk and return of all the MV optimized portfolios with estimation error. The solid blue curve represents the out-of-sample average risk and return of all the REF optimized portfolios with estimation error.

Results

The dashed and solid blue curves represent the in-sample and out-of-sample REFs; in other words what you use to invest versus what happens on average in the investment period. These curves intersect. Given the congruity between forecasts and actual performance on average, we can conclude that the inputs are very useful, and that RE optimization uses the information very well on average.

In contrast, the dashed and solid red curves representing the in-sample and out-of-sample MV efficient frontiers are far apart. The in-sample MV efficient frontiers overestimate the return associated with portfolio optimization not only with respect to RE (blue dashed curve) but importantly with respect to out-of-sample investment performance (red solid curve). Even with very useful inputs, MV efficiency maximizes the errors in the risk and return inputs, creates upward-biased estimates of future performance, and substantially underperforms RE optimization on average. The same reliable investment information that performed so well with REF portfolios is misused by MV efficiency. In addition, the error maximization property of MV efficiency means that real estimates such as trading costs are misused in the optimization process; you are likely to think that the returns are much higher relative to trading costs than they actually are.[vii]

These results should trouble many MV efficiency investors and advisors. Resource allocation bias toward investment forecasting and away from effective optimization technology may often be self-defeating. While good inputs are important, they need to be transformed into optimized portfolios that do not misuse the information.

We have demonstrated that RE optimization is a necessary condition for effective portfolio optimization. The resampling process unbiases the optimization process so that the information is transferred more directly into the optimized portfolios.

It also should be clear that RE optimality is not inconsistent with any input optimization process that may reduce estimation error. The better the input estimates in RE optimization the more likely investment performance is improved. In particular, various statistical estimation techniques in conjunction with REF technology holds out the promise of dramatically improved optimized portfolio performance.

Understanding Resampled Efficiency Optimality

Unlike MV efficiency, RE optimization is an out-of-sample definition of portfolio optimality. MV efficiency is correct only when there is no uncertainty in the optimization inputs. But investment forecasts always are uncertain. RE optimization deals with uncertainty by simulating all the many ways markets and assets can perform based on your forecasts and then finding portfolios that, on average, do well with respect to all the simulated outcomes.

Portfolio optimization with uncertainty implies fundamental changes in investor perceptions. The MV efficient frontier familiar to students of finance and investment professionals turns out to be essentially useless in understanding portfolio optimality. In particular, where a portfolio plots in the mean-variance diagram may not determine even relatively to others how it is likely to perform.

Note that the REF is shorter than the MV frontier in figures 1 and 2. Is this a symptom of something amiss with our definition of efficiency? Are there portfolios that are more optimal than RE optimization? The paradox is easily explained. If you are 100% certain of your risk-return estimates, then Markowitz efficiency is for you. If you are less than 100% certain, you expect less return and are less willing to put as much money at risk. That is why the REF is shorter and below the classical Markowitz frontier. To drive the point home, suppose you are 100% uncertain of your information. In this case, the REF portfolios should represent no information, and optimality is either the benchmark or equally weighed portfolio. The REF, in this extreme uncertainty, collapses to a point. RE efficiency is different, because it takes uncertainty into account while defining optimality. RE optimization is the paradigm of choice for defining optimized portfolios under the condition of uncertain information.

Figure 2 shows the usual in-sample relationship between the MV and REF. Now consider a portfolio that plots at point A above the REF and below the MV frontier. Is portfolio A more efficient than the RE optimized portfolios? Do you prefer investment in portfolio A to a portfolio with similar risk on the REF?

FIGURE 2: MEAN-VARIANCE AND RESAMPLED EFFICIENT FRONTIERS

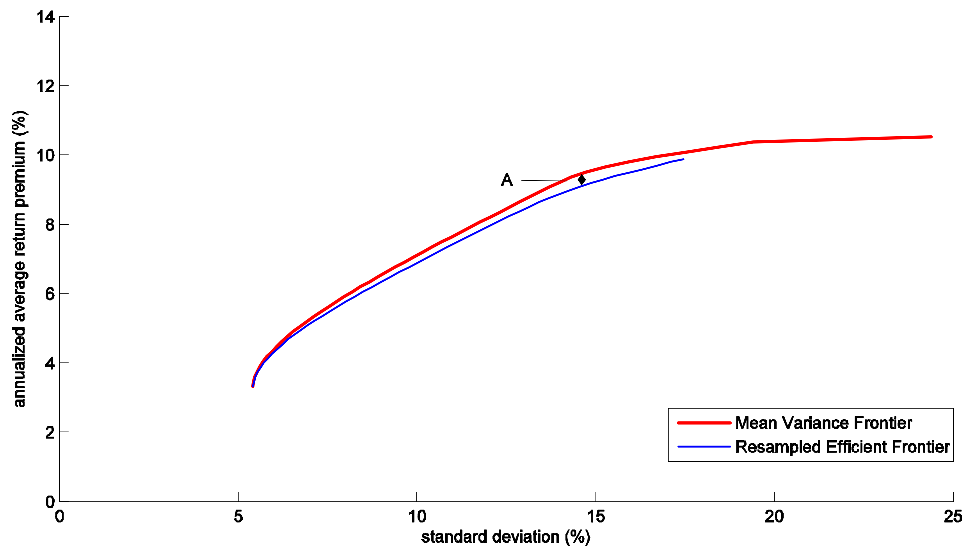

By definition, portfolios on the REF are optimal conditional on the amount of uncertainty in your information. The portfolio that plots at A is not preferable. Intuitively, the asset weights may be "too active" relative to the level of information in your inputs. But an additional point helps to further clarify the issue. A portfolio that plots at point A is not unique. An infinite number of portfolios have the same mean and variance as A. One-asset portfolios may even exist that plot at A yet have much risk out-of-sample and are clearly inefficient by anyone's definition. This discussion highlights the fact that where a portfolio plots in the MV graph may say very little about whether it is a good investment.

The underlying financial reality explained by RE optimization is that the structure of the portfolio, not its mean and variance parameters, defines an investment useful optimality.[viii] This is what has been missing in our understanding of portfolio efficiency for nearly 50 years.

Conclusion
Portfolio structure relative to out-of-sample performance conditional on forecast certainty characterizes a more useful definition of portfolio optimality in investment practice. Portfolio optimality defined by out-of-sample investment performance reveals many investment illusions that negatively affect investment practice.[ix] RE optimization offers important new investment tools and more effective and intuitive asset management.

Endnotes

[i] Jobson and Korkie (1981) use a simulation study framework to prove these results. In a more recent study, DeMiguel et al. (2007) test the performance of 14 estimation models of unbounded MV portfolio choice in the context of estimation error and find, as in Jobson and Korkie (1981), that none seem to be reliable improvements over equal weighting.

[ii] The Resampled Efficient Frontier technology is protected by U.S. and foreign patents and patents pending worldwide. It first was proposed in Michaud (1998) and recent updates in Michaud and Michaud (2008a: 2008b). New Frontier Advisors is exclusive worldwide licensee.

[iii] More details of simulation tests are given in Michaud (1998), Michaud and Michaud (2008a: 2008b), and Markowitz and Usmen (2003).

[iv] The data in this example are taken from chapter 2 of Michaud (1998).

[v] We use two Stein estimation procedures: James-Stein-Efron-Morris for return and Ledoit for covariance estimation. See Michaud (1998: 2008a) for further information and references.

[vi] In these tests the Stein estimates of optimization inputs are surrogates for reliable forecasts of risk and return and are not meant to replace the process of developing reliable investment information from economic, market, and other sources.

[vii] A very sophisticated study, addressing the same issues, was performed by Markowitz and Usmen (2003). They used a more sophisticated Bayesian approach for defining risk-return estimates. Their study compared the unenhanced risk-return estimates and RE optimized portfolios to the Bayesian enhanced estimates with Markowitz optimization. Their results are consistent with those here. They found that RE optimized portfolios outperformed MV optimization on average and in every one of 30 individual tests even with inferior risk-return estimates.

[viii] For example, rebalancing or other procedures based solely on a portfolio's mean and variance parameters are unlikely to have useful investment value out-of-sample. There are a number of fundamentally important associated issues of misunderstandings of portfolio optimality that are the consequence of ignoring estimation error but beyond the scope of this report. See Michaud and Michaud (2005) for further discussion.

[ix] See Michaud and Michaud (2008a; 2008b) for further discussion.

References

DeMiguel, V., L. Garlappi, and R. Uppal. 2007. "Optimal Versus Naïve Diversification: How Inefficient is the 1/N Portfolio Strategy?" *Review of Financial Studies*. Advance Access published December 3, 2007, doi: 10.1093/rfs/hhm075.

Jobson, J.D., and Bob Korkie. 1980. "Estimation for Markowitz Efficient Portfolios." *Journal of the American Statistical Association* 75 (371): 544–555.

——— 1981. "Putting Markowitz Theory to Work." *Journal of Portfolio Management* 7 (4): 70–74.

Markowitz, H. 1959. *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley.

Markowitz, H., and N. Usmen. 2003. "Resampled Frontiers versus Diffuse Bayes: An Experiment." *Journal of Investment Management* 1(4): 9–25.

Michaud, R. 1989. "The Markowitz Optimization Enigma: Is Optimization Optimal?" *Financial Analysts Journal* 45 (1): 31–42.

———. 1998. *Efficient Asset Management*. New York: Oxford University Press.

Michaud, R. and R. Michaud. 2005. "The Fundamental Law of Mismanagement." New Frontier Advisors Newsletter, July. www.newfrontieradvisors.com/about/research/Articles.html.

______. 2008a. *Efficient Asset Management*. 2nd ed. New York: Oxford University Press.

———. 2008b. "Estimation Error and Portfolio Optimization: A Resampling Approach." *Journal of Investment Management* (first quarter): 8-28.