Book Review

Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation. Richard O. Michaud. Boston, MA: Harvard Business School Press, 1998. 152 pp. ISBN 0-87584-743-9.

Markowitz's portfolio theory is one of the most important theoretical developments in finance. Its elegance and theoretical appeal are illustrated by the fact that it is taught in virtually every undergraduate and MBA level course on portfolio management. Yet the Markowitz mean-variance (MV) optimization is not widely used by the investment community. In this book, Richard Michaud, a leading researcher in this area, explains why this is so. Based on his diagnosis of the problem, he advocates a "statistical view" of the MV optimization that leads to new procedures that can reduce or eliminate many practical deficiencies of the MV optimization.

Using a simple global asset allocation example throughout the book, he illustrates the problems associated with a naïve implementation of the MV optimization procedure. He then describes how the new procedures can be applied, and shows how they improve the MV optimization results. The bottom line: the "statistical view" and the new procedures offer the potential of substantial improvements in the investment performance of the MV optimization. I agree with this assessment. I find this book a must-read for industry professionals and researchers who are interested in this area.

Chapters 1–3 introduce the MV efficiency framework, the traditional criticisms about the framework, and the commonly suggested alternatives. While Chapters 1–3 should be familiar to an average MBA student, Chapter 4 illustrates the practical problems of the MV optimization that are probably not familiar to someone who has not tried it. The problem is that the optimal portfolio is highly sensitive to small changes in the optimization input, and out of the sample it performs very poorly on average (much worse than the equal-weighted portfolio) in realistic situations. Anyone familiar with Matlab (SAS Institute, Cary, NC) can conduct the following simple simulation in a matter of 15 minutes to see this is the case: (1) You assume a realistic expected return vector and the return covariance matrix. (2) Based on this, you simulate 20 years of monthly returns. (3) Using the simulated returns, estimate the mean returns and return covariance matrix, and use them to find the MV-efficient portfolios. (4) Then check the performance of these portfolios using the true expected return vector and return covariance matrix. If you repeat steps (2)–(4) many times, you can see that there is an enormous variability of the portfolios' performance, and the average performance is quite bad. You will also see high dispersions in the portfolio weights among those efficient portfolios that theoretically should be identical. This illustrates the major problems of a naïve implementation of the MV optimization.

The rest of the book discusses various statistical techniques that can enhance the performance of the MV optimization. Chapter 5 discusses how to test whether a given portfolio is MV efficient. Given the wide variability in the efficient portfolios (either their performance or their portfolio weights) that we already observed in Chapter 4, it is quite clear that such tests cannot have much power. Therefore, I find this chapter of limited value.

Chapter 6 introduces a new concept of efficiency, the "resampled efficient frontier." In the above simulation, we can check the average performance of the efficient portfolios. We can also check the performance of the "averaged" efficient portfolios, averaged over different simulations. Michaud called the properly averaged efficient portfolios the resampled efficient frontier. The resampled efficient frontier might perform better than the traditional one since it averages the portfolio weights over many simulations, so its weights are more stable. Nevertheless, we cannot expect it to do much better in practice. This is because, in practice, we do not know the true mean return vector and return covariance matrix. We can only estimate them, using observed data, and then construct the resampled efficient frontier using these estimates. The resampled efficient frontier might perform well if evaluated using the estimated mean return vector and the estimated covariance matrix. However, the whole issue is that these estimates are very noisy. So any portfolio derived from them without using extra-sample information will also be very noisy and cannot expect to perform well if evaluated using the true expected returns and return covariance matrix. Specifically, to the extent that we are only averaging over the portfolios that are all derived from the same single estimate of the mean returns and return covariance matrix, and thus all the portfolios inherit the deviation of this estimate from the true one, the deviation of these portfolios from the true efficient frontier is all highly correlated, hence averaging will not help much. Indeed, Michaud only finds marginal improvement in performance of the resampled efficient frontier over the traditional efficient frontier.

Chapter 7 talks about portfolio revision and confidence regions. Given the expected return vector and the return covariance matrix, we can simulate many sets of returns and construct the efficient frontier from each set of simulated data. In effect, we have a distribution of the efficient frontiers. Based on this, we can conduct statistical tests. I find this is much the same topic as the one in Chapter 5. I would prefer to combine the two chapters. On the value of such tests, again I would say that, given the huge variability in the simulated efficient frontiers that we already observed in the previous chapters, such tests clearly cannot have much power.

Chapter 8 discusses some improved estimations of the expected returns and return covariance matrix. I think finding better estimates is the most promising way to enhance the practical performance of the MV optimization. This chapter discusses the James–Stein estimator of the mean returns, the Frost and Savarino estimator that jointly estimates the mean returns and return covariance matrix, and Ledoit's shrinkage estimator of the covariance matrix. All these estimators are Bayesian or at least have Bayesian interpretation. This is no surprise. The root of the practical problems of the MV optimization is that estimates that only use sample information are too noisy. The only way to overcome this problem is to incorporate extra-sample information, and the Bayesian method is the natural approach to do that. In concluding this chapter, Michaud states that "the development of Stein estimators for portfolio optimization is at a relatively early state. In particular, the identification of the optimal Stein estimator in many practical contexts is generally open." This area definitely offers the greatest potential to enhance the MV optimization performance and deserves further research.

Chapter 9 talks about benchmark active asset allocation. This is just the MV optimization applied to the returns in excess of a benchmark portfolio. So all the discussions in the previous chapters also apply to this case. Chapter 10 is about investment policy and economic liabilities, which is important for managers of pension funds and endowments. Chapter 11 deals with how to combine exogenous return forecasts with historical data in estimating mean returns and return covariances. I believe that the most natural way to do this is the Bayesian approach, treating the exogenous return forecasts as the prior, and using historical data to update it to arrive at the final estimate. Surprisingly, this is not the approach advocated in the book. Instead, it advocates a Theil and Goldberger regression approach. In my view, by following a Bayesian approach, we can tie Chapter 11 and Chapter 8 together under a unified theme. The theme is to incorporate extra-sample information to improve the estimates of the optimization inputs.

Chapter 12 outlines ways to avoid optimization errors. Obviously this is very important to industry professionals.

In summary, this book not only is highly valuable to industry professionals, but also introduces academics to many research questions. Michaud gives an excellent discussion regarding the research questions at the end of Chapter 8, that is, how to identify the optimal Stein estimator to jointly estimate the mean returns and return covariance matrix? In Frost and Savarino's estimator, the prior assumes equal expected returns, return covariances, and return variances. Previous research [Green and Hollifield (1992)] shows that this may not be a good prior. For U.S. domestic stocks, a better one would be a one-factor structure, as in Ledoit's estimator. Accepting a one-factor structure and the APT would then tie the estimates of the mean returns and return covariance matrix together. MacKinlay and Pastor (2000) and Wang (2001) exploit such a relationship to improve the return estimates.

Researchers also find that imposing upper and lower bounds on portfolio weights can improve the MV optimization results [Frost and Savarino (1988)]. This issue is not discussed in this book. The practice of imposing bounds on portfolio weights can be considered as incorporating extra-sample information (or belief). Recent work by Jagannathan and Ma (2000) provides some insights as to why imposing such bounds may help.

Tongshu Ma
University of Utah

References

Frost, P. A., and J. E. Savarino, 1988, "For Better Performance: Constrain Portfolio Weights," *Journal of Portfolio Management*, 5, 29–34.

Green, R. C., and B. Hollifield, 1992, "When Will Mean-Variance Efficient Portfolios be Well Diversified?" *Journal of Finance*, 47, 1785–1809.

Jagannathan, R., and T. Ma, 2000, "Covariance Matrix Estimation: Myth and Reality," working paper, Northwestern University and the University of Utah.

MacKinlay, A. C., and L. Pastor, 2000, "Asset Pricing Models: Implications for Expected Returns and Portfolio Selection," *Review of Financial Studies*, 13, 883–916.

Wang, Z., 2001, "A Shrinkage Theory of Asset Allocation Using Asset-Pricing Models," working paper, Columbia University.