Portfolio Optimization, Heuristics, and the "Butterfly Effect"

David N. Nawrocki, Villanova University

Introduction

The "Butterfly Effect" has been an interesting concept in the development of chaos theory. Given a complex system like the global weather system, the butterfly effect is rather intriguing. It says that the flapping of a butterfly's wings in Beijing will work its way through the system and result in a tornado in Oklahoma.

The same effect is at work with portfolio optimizers that perform asset allocation and portfolio allocation chores. A small change to an input works its way through the system of equations and results in a large change in allocations. As a result, it is very easy to arrive at a set of "non-intuitive" allocations; i.e. they don't exhibit common sense. Once a set of portfolio allocations is put into place, a small change in the market will result in large changes (sometimes negative) in the portfolio returns. Practitioners, understandingly, have been losing confidence in the results of portfolio optimizers. The purpose of this paper is to explore the reasons why and to offer one reasonable alternative, a portfolio heuristic.

A ***portfolio optimizer*** is a solution algorithm used to determine the ***optimal or the very best solution*** given a set of information inputs. A ***portfolio heuristic*** is a solution algorithm used to determine ***"an approximately good" solution*** given the same set of information inputs. A heuristic does not provide an optimal solution, but it does provide a reasonably good solution. Its main advantage is that a heuristic is cheaper, faster to use than an optimizer, and it is less sensitive to the "butterfly effect". In this paper, the basic inter-workings of a heuristic algorithm will be presented.

Statistical Problems with Optimizers

Hensel and Turner (1992) identify four problems associated with the "butterfly effect" and portfolio optimizers and suggest solutions to these problems:

1. Inappropriate allocations – Securities with extreme values for returns and variances will be over weighted or under weighted in the portfolio. Securities with large returns and low variances will be over weighted. Securities with low returns and high variances will be under weighted. Therefore, if there are any errors in estimating the returns and variances for the securities, there will be over weighted or under weighted allocations. The likelihood of making estimation errors is very high. Therefore, it is very likely that the final portfolio will contain a heavy weight into one asset that is counter intuitive to a portfolio manager. (See Michaud, 1989).

2. Instability of the optimal solution -- Small changes in inputs, especially mean returns, can cause large changes in the optimal asset weightings. Therefore, optimal weights

will change significantly over time. As the manager makes new estimates for risk and return and revises the portfolio, the optimal weights will be excessively unstable. The instability of the optimizer weights over time is the direct result of making estimation errors. (See Kallberg and Zeimba, 1984 and Adler, 1987).

3. Inclusion of inappropriate assets -- Hensel and Turner (1992) note that the inclusion of any non-market-traded asset types is inappropriate. Real estate is one such investment asset. Any statistical measure of risk is not going to capture the illiquidity of real estate or the appraisal bias of returns. Also, transaction costs are significantly higher in real estate than in other security markets. Most optimizers will over weight the portfolio allocation into real estate because the statistical risk measure does not capture the real risk of an illiquid market.

4. Non-uniqueness of optimizer solutions -- Because estimation error prevents the determination of the one set of estimated returns and variances, there isn't just one solution from the optimizer. There is one optimal solution; however, the optimizer isn't finding it. Given different errors in the inputs, different alternative solutions that are non-optimal will be returned. The analyst will have no idea which solution is optimal and which one is not. (See Michaud, 1989).

Inappropriate Use of Optimizers by Practitioners

First, it should be noted that this is not the fault of practitioners. They have been sold an academic tool not a practitioner tool. In addition, there is a problem with the ***inappropriate use of optimizers***. The most important use of optimizers by practitioners is to generate a portfolio that will provide acceptable performance ***in the future,*** in other words, forecasting. A major problem of using optimizers to forecast is that optimizers will overfit the past. This results from academics trying to explain how systems work by studying the past. Academic models will overfit the past data in order to provide greater explanatory power. Unfortunately, this approach is rarely a very good forecaster and provides excessive amount of estimation errors. A number of academic models have successfully explained the past but have proven to be terrible forecasters.

Overfitting the past data may explain 99% of what happens during the past, but may only forecast 5% of what happens in the future. Reducing the complexity of the model may reduce the explanatory power to 60-70% but forecasting ability may increase to 40-50%. The better the forecasting ability the smaller the estimation error.

```
               Past                    Present            Future
<----------------------------------------------|------------------------------------------->
    Academics Explain                          Practitioners Forecast
    Academics Look Backwards                   Practitioners Look Forward
    Optimizers                                 Heuristics
    Complexity                                 KISS (Keep It Simple, Stupid)
```

Academic models, which include optimizers, are backward looking models. Practitioners really don't have as much need for optimizers as their interest is not in explaining the past. They need to use forward-looking models to minimize the “butterfly effect”.

Possible Solutions

Hensel and Turner (1992) identify two approaches to (1) reduce the impact of estimation errors, (2) lead to more diversified portfolio weights, and (3) provide stable portfolios that change allocations slowly over time. One approach is to methodically massage the inputs and the second approach is to constrain the optimal weights.

Jobson and Korkie (1981) and Michaud (1989) suggest Stein estimation (shrinkage towards single mean return for an overall asset class) as one method of massaging the inputs. Hensel and Turner (1992) cover a number of studies that have used Stein estimators for optimizer inputs. Typically, these studies concentrate on the mean returns as these seem to be the input most sensitive to the "Butterfly Effect". One example would be to substitute a general asset class return for the returns of the individual assets. Specifically, the mean return of an industry can be substituted for the mean returns of the individual stocks. Alternatively, the return for a broad market index such as the S&P 500 would be used for large cap growth, mid-cap value, small-cap growth and so forth.

Elton, Gruber, and Padberg (1976) and Elton, Gruber and Urich (1978) developed a successful portfolio heuristic by using a single average correlation coefficient. A 20 asset optimization problem will contain 190 correlation coefficients. Elton, et al substitute one correlation value (the average) for the 190 correlation values. They demonstrate that this heuristic approach will forecast better, provides stable portfolio allocations, and provides more diversification than an optimizer. Therefore, Stein estimation can help develop portfolio heuristics as well as change the inputs into a portfolio optimizer.

The second approach is to add a maximum allocation constraint for each asset when running the optimizer. For example, each asset class may be constrained to a maximum of 15%. Frost and Savarino (1988) tested this approach for portfolios of individual assets while Chopra (1991) tested this approach for general asset classes.

Problems with Optimizers

However, there are problems using constraints with optimizers. The major problem is a degrees of freedom problem. If the number of assets in the optimization problem exceeds the number of observations used to estimate the optimizer's inputs, then there is essentially negative degrees of freedom. Adding maximum allocation constraints to the optimization problem will effectively double or triple the number of assets and further reduce the degrees of freedom.

As there are many ways to do a job, there are many ways to find an optimal set of allocations for a portfolio. The steps that are computed to find an optimal solution are called an ***algorithm***. There are many optimizer algorithms, some of which do a better job than other algorithms.

Markowitz (1987) deals with this problem by studying how sensitive different optimal algorithms are to negative degrees of freedom. Adding constraints to the problem doubles or triples the number of assets in the optimization problem. Effectively, adding a minimum investment or maximum investment constraint is equivalent to adding a new security to the problem. There are three general types of optimization algorithms. All three of the following will perform a mean-variance optimization.

1. Calculus minimization--simultaneous equations optimizing algorithm. This approach is probably the first portfolio optimizer, first described by Martin (1954). This algorithm is covered extensively in portfolio management textbooks like Francis (1991) and Elton and Gruber (1995). The major drawback is that this algorithm can only handle up to 20-30 assets before it runs into input estimation errors. (See Nawrocki, 1996) Therefore, a 10-asset optimization problem with 10 maximum allocation constraints is going to be sensitive to estimation errors and to the number of observations because effectively there are 20 assets. It is not a stable algorithm and highly sensitive to the “butterfly effect”. As a result, many times this algorithm will not even return an optimal portfolio.

2. General nonlinear programming optimizing algorithms. The Markowitz optimization problem is a problem in nonlinear optimization. However, there are many other nonlinear optimization problems to be solved, therefore, a general algorithm is required. The nonlinear programming optimizers in programs such as EXCEL are examples of this algorithm. These algorithms are going to be sensitive to the number of observations problem because of their complexity. The complexity derives from the very nature of a general algorithm that has to be able to solve any nonlinear programming problem not just a portfolio optimization problem. Another problem is that nonlinear programming algorithms have to be supplied with a first guess as to the final solution. In other words, the user has to guess at the initial set of portfolio weights. If the first approximation is accurate, then the algorithm will converge on the overall optimal solution. If the first approximation is not accurate, the algorithm may converge on a non-optimal solution rather than the optimal solution. This chance of converging to a non-optimum solution is a real source of problems with this algorithm. EXCEL is a particular problem because the user has to know to provide the original set of allocations and to keep solving the problem iteratively until the solution stabilizes on a particular set of allocations. This algorithm will also break down rapidly with a small number of assets and constraints.

3. Critical line optimizing algorithm. -- Markowitz (1959) originally developed this algorithm to specifically solve the portfolio optimization problem. It does not have the complexity of the general nonlinear programming optimizing algorithms and therefore, it shows very little sensitivity to the number of observations no matter how

many assets are optimized (Nawrocki, 1996). Therefore, this is the best algorithm for adding constraints to the optimizer. The main drawback to this algorithm is that it does not solve for specific points on the efficient frontier, but rather provides a sample of the portfolios on the efficient frontier. It provides a portfolio everyplace on the frontier where an asset enters or leaves the optimal portfolio. It does not provide a portfolio for a specified return. Therefore, very few commercial packages use this algorithm which is unfortunate as it is the least sensitive to the "butterfly effect".

Even when using the Markowitz critical line optimal algorithm, maximum allocation constraints force the optimizer so far away from the optimal solution that the final solution is one of many non-optimal solutions. Since a non-optimal solution is being derived, the optimizer does not have any inherent advantage over a heuristic non-optimal solution. There is no advantage plus we have the additional computational complexity and no promise of a better forecast.

Heuristics Algorithms as a Solution

Sharpe (1967) seems to be the first academic researcher to abandon an optimizer when faced with maximum allocation constraints. He was consulting with a mutual fund that could not have more than 5% of its assets in any one stock. His solution was to use a heuristic that he developed. First, he computed a risk-adjusted return for each stock using the beta as a risk measure. Then, he ranked the stocks from the highest risk adjusted return to the lowest risk adjusted return. He then invested 5% of the portfolio in the top 20 stocks. In this way, he met the 5% constraint. Having 20 stocks in the portfolio is also important since beta only measures nondiversifiable market risk. It ignores the diversifiable risks that make a company unique. By including 20 stocks in the portfolio, Sharpe hopefully has enough stocks to diversify away the diversifiable risk leaving only the market risk inherent in the stocks to affect the portfolio. The beta risk measure has proven to be a pretty good forecast of the future as long as a portfolio consists of 20 or more stocks. Elton, Gruber and Urich (1978) studied the forecasting abilities of the beta and found it was superior to using a full correlation matrix to forecast a future correlation matrix. [They also found that the average correlation coefficient used in the Elton, Gruber and Padberg(1976) was the best predictor of a future correlation matrix.]

Nawrocki(1983, 1990) has demonstrated that a simple risk-reward heuristic derived from the Elton, Gruber and Padberg(1976) heuristic using the downside risk measure (lower partial moment, or LPM) will provide better investment performance than an optimizer over 30 year investment horizons. Computing portfolio allocations with this algorithm is quite easy. First, the portfolio manager decides how many assets will be in the portfolio. The reward to LPM ratio (reward to semivariability when the LPM degree $n = 2.0$) is computed for each asset using the following formula: (Security Return – T-Bill return)/Security Semideviation. The portfolio weights are determined by adding the ratios for the number of assets in the portfolio and dividing the sum into each of the individual R/LPM ratios.

An Example of Optimizers and a Heuristic

The best way to demonstrate the inter-workings of the optimizers and a heuristic is through an example. ***The reader should be cautioned that an example is not an empirical test demonstrating the superiority of one technique over another, but rather an example of what might be expected from the techniques***. Over a short holding period as used in this example, one algorithm will not always dominate the other algorithms. However, the results presented in this example are consistent with long run results obtained over 20-30 year time horizons by previous studies (Nawrocki, 1983, 1990).

The example is presented in three exhibits and utilizes monthly total returns for the 35 Fidelity Select Sector Funds obtained from the Morningstar Principia Plus database. The data were obtained from April 1994 to March 1998. The portfolios were generated and their performance was evaluated using the PMSP Professional optimization program.

First, three portfolios are generated using data from April 1994 to March 1997. This constitutes a historic period. Out-of-sample results were obtained from a holding period from April 1997 to March 1998.

The three portfolios were generated using three different portfolio algorithms. The first portfolio is a *mean-variance optimization* generated using the Markowitz quadratic programming critical line algorithm with a variance-covariance matrix. The second portfolio is a *mean-semivariance optimization* generated using the Markowitz critical line algorithm with a semivariance-semicovariance matrix. The ***semivariance*** is sometimes referred in the literature as the ***downside risk measure*** or the ***downside deviation measure.*** (See Nawrocki, 1999).

The first two algorithms are optimizers, the only difference being the risk measure. One uses the variance and the second uses the semivariance. The third algorithm is the *risk-reward heuristic* using the reward-to-semivariability ratio (R/SV Ratio) to compute portfolio allocations. As a performance benchmark, an equal weight portfolio consisting of all 35 Fidelity Select Funds is also evaluated.

The following results from the historic period should be noted in Exhibit A.

1. The *risk-reward heuristic* was generated using 10 assets in order to constrain the maximum allocation to some extent. The allocations for each fund are very easy to compute. First, the R/SV ratio for each fund is computed and sorted from the highest ratio value to the lowest ratio value. The top 10 funds are summed (7.6595). This sum is divided into the R/SV ratio for each fund. Therefore, the allocation for Health Care is computed by dividing 1.3762 by 7.6595 which yields the allocation of 17.9677%. If we were attempting to reduce the maximum allocation further, we would simply increase the number of funds in the portfolio.

2. Right away, the disadvantage of the *mean-variance optimizer* and the *mean-semivariance optimizer* is apparent. Both portfolios are essentially three fund portfolios. One fund, Food & Agriculture, dominates both portfolios with over 46% and 58% of the portfolio allocation. The *risk-reward heuristic*, on the other hand, distributes its allocations among 10 funds with the largest allocation to Health Care of 17.97%. The allocations are distributed fairly smoothly down to the 5.90% allocation to Financial Services.

Exhibit A - Portfolio Allocations Using Historic Period April 1994 to March 1997

Fidelity Funds	Mean-Variance Optimizer	Mean-Semivariance Optimizer	Risk-Return Heuristic	R/SV Ratio
Health Care	20.2406	20.8332	17.9677	1.3762
Food & Agriculture	46.0783	58.4028	13.4492	1.0301
Insurance			10.1752	0.7794
Home Finance	6.0009		10.1708	0.7790
Regional Banks			10.0225	0.7677
Defense & Aerospace	0.9643		8.4518	0.6474
Energy Services	21.7680	12.0244	8.1984	0.6280
Electronics			8.0149	0.6139
Financial Services			7.6471	0.5857
Industrial Equip.			5.9025	0.4521
Growth Utilities	1.7854			0.4001
Paper & Forest	3.1543	4.8444		0.2885
Transportation		3.8953		0.1939
Sum of Allocations	100.0000	100.0000	100.0000	
Sum of Top 10 R/SV Ratios				7.6595

R/SV Ratio = (Return – T-Bill Rate)/Semideviation

Source: Morningstar Principia Plus for Mutual Funds

Exhibit B provides the summary statistics of the three portfolios during the historic period from April1994 to March 1997. The results indicate that the three portfolios from three different algorithms are comparable in historic performance.

1. The *risk-return heuristic* portfolio is not optimal. While it has a higher return than the two optimizers, it has higher standard deviation and semi-deviation values than the two optimizer portfolios. It also has lower R/V and R/SV ratios than the optimizers but, remember that the purpose of the heuristic is to forecast not explain.

2. The optimizer results are the expected results. The *mean-variance optimizer* optimizes the portfolio variance so it enjoys a lower standard deviation than the *mean-semivariance optimizer*. The *mean-semivariance optimizer*, on the other hand, optimizes the portfolio semivariance so it has a lower semi-deviation (square root of the semivariance) than the *mean-variance optimizer*.

3. The critical line algorithm does not optimize for a specific return. It provides a sampling of the efficient frontier. The two optimizer portfolios are comparable on a risk-reward basis. Their variances are close to each other and their R/SV ratios are comparable.

Exhibit B - Historic Performance for the Three Exhibit A Portfolios from April 1994 to March 1997

	Mean-Variance Optimizer	Mean-SemiVariance Optimizer	Risk-Return Heuristic
Annualized Return	25.2686	23.2620	**26.4002**
Monthly Mean Return	1.8952	1.7581	**1.9715**
Standard Deviation	**2.1589**	2.1669	2.7428
Semi-Deviation	0.8947	**0.8180**	1.2981
R/V Ratio	**0.6891**	0.6233	0.5703
R/SV Ratio	**1.6629**	1.6513	1.2050

Source: Morningstar Principia Plus for Mutual Funds
Bold Print indicates best performance of the three algorithms.
Semi-Deviation is the square root of the semivariance.
R/V Ratio = (Return – T-Bill Rate)/Standard Deviation
R/SV Ratio = (Return – T-Bill Rate)/Semi-Deviation

Using the holding period of April 1997 to March 1998, Exhibit C provides an idea of how a heuristic might stack up against an optimizer in performance. Over a short time horizon, the heuristic will not always outperform an optimizer, however, in every long-term study, it has outperformed optimizers. The performance in Exhibit C is typical of the long-term performance of the three algorithms.

1. The *mean-semivariance optimizer* portfolio is the most concentrated portfolio (5 funds) and it seems to have performed the worse. It has the lowest return and lowest R/SV ratio. It does have lower standard deviation and semi-deviation values than the *mean-variance optimizer* and the *risk-reward heuristic* portfolios. Therefore, it does manage a higher R/V ratio than the *mean-variance optimizer* portfolio. The *mean-variance optimizer* portfolio has a few more funds in it (7). It has the second highest

return but with higher standard deviation and semi-deviation, its risk-reward performance relative to the *mean-semivariance optimizer* is mixed.

2. The *risk-reward heuristic* has the highest return. It also has the highest standard deviation but it exhibits the lowest semideviation of the three portfolios. Therefore, it did the best job of forecasting the future downside risk. This is demonstrated by the highest R/SV ratio in the holding period. It also had the highest R/V ratio. Overall, the best forecasting performance was turned in by the *risk-reward heuristic*.
3. The Equal Weight benchmark was outperformed on a risk-return basis by all three algorithms. Therefore, each algorithm provided improvement in forecasting ability over a naive portfolio strategy of investing an equal allocation into each of the 35 Fidelity Select funds.

Exhibit C – Holding Period Performance for the Three Exhibit A Portfolios from April 1997 to March 1998

	Mean-Variance Optimizer	Mean-SemiVariance Optimizer	Risk-Return Heuristic	Equal Weight
Annualized Return	40.8538	37.5325	**47.8451**	40.5205
Monthly Mean Return	2.8957	2.6913	**3.3120**	2.8754
Standard Deviation	3.7465	**3.3080**	4.0627	4.1594
Semi-Deviation	1.3855	1.3488	**1.3451**	1.6327
R/V Ratio	0.6642	0.6904	**0.7149**	0.5934
R/SV Ratio	1.7959	1.6933	**2.1594**	1.5116

Source: Morningstar Principia Plus for Mutual Funds
Bold Print indicates best performance of the three algorithms.
Semi-Deviation is the square root of the semivariance.

Implications for Practitioners

Portfolio optimizers are one of the great garbage-in garbage-out information processors in the financial field. Practitioners have to exercise the utmost care about what numbers they input to an optimizer. Dumping 20 years of monthly data on broad asset classes into an optimizer will earn you what you deserve: poor performance.

Portfolio optimization was originally developed to diversify stock portfolios not broad asset classes. The reason portfolio optimization works with stock portfolios is that the correlation coefficients between stocks are close to zero. However, with asset classes and mutual funds running correlation coefficients over 0.70, the optimizer does not add a great deal of value and introduces unneeded complexity to a sensitive forecasting system. The high correlation coefficients for most asset classes are one reason why the real estate indexes were popular in optimization solutions. They were the only asset class with a lower correlation with other asset classes. Unfortunately, it is very difficult to obtain the return and liquidity promised by the real estate indexes. Investing in an artificial index that does not reflect real world market conditions will also earn well-deserved poor performance.

The key to a successful portfolio allocation decision is to have very good estimates for risk and return. The makeup of the portfolio can be determined heuristically through risk-return ratios at the general asset class level, at the mutual fund level and at the individual stock level.

To comment on the Jahnke (1997) asset allocation hoax controversy, running an optimizer to determine asset allocation (strategic or tactical) by itself does not add value to a portfolio. It is the careful selection of assets and the careful determination of risk and return measures that the manager inputs to the optimization or the heuristic algorithm decision that provides the added value.

A Personal Odyssey

I wrote my first portfolio optimization computer program using the simultaneous-equations algorithm for my senior year undergraduate project almost thirty years ago. It was a crude affair, like the first telephone, but it optimized and generated an efficient frontier for 10 stocks. It was a great deal of work because it could only optimize one portfolio at a time. So for each portfolio that I generated on the efficient frontier, I had to hand a card deck to a dispatcher. I would receive my printout and card deck back the next day and I could add another point to my efficient frontier graph. (As an undergraduate, I had the lowest priority for the running of my programs on the university main frame.)

I worked in a bank trust department in the early 1970s as an economic analyst for the security analysts and the portfolio managers. The trust department was organized specifically to the traditional portfolio management process. I quickly found out that portfolio managers were not interested in my little toy because it was completely foreign to how they managed portfolios. Even if they were interested, portfolio optimization could not have been implemented there because the organizational structure was incompatible with this new portfolio management process. So I learned not to try to change people or organizations when they did not want to be changed. The advantage that a small investment management practice has over a larger investment firm is that it is easier to change the organizational structure of the practice to encompass the modern portfolio theory approach to portfolio management. I would venture to say that the big failure of modern portfolio theory in the 1970s was due to organizational structure problems in the large organizations that could afford a main frame computer rather than a failure of the theory.

I continued my education and became a finance professor teaching portfolio theory using my software. In 1982, I started to market my software to universities. My PMSP package contained a number of optimizers, heuristic algorithms, and downside risk measures. I also started to make sales to large investment firms and banks. My typical commercial client had a Ph.D. in some field or was an MBA/CFA who had learned portfolio theory in school. Almost everybody bought my package for the optimizer and almost everybody spent a month on the phone with me learning how to use it. As some of my customers became more familiar with the actual workings of an optimizer and asked me how I would approach their problem, I would slowly introduce them to the heuristic algorithms and the downside risk measures. My client feedback may be biased, but I noted that my most successful customers who stayed with the software were the ones who used the heuristics and the downside risk measures.

In 1992, I finally ventured out of the academic ivory tower and started to learn portfolio management in the real world. Working with a former student, we did a lot of backtesting of portfolio strategies, working with a small investment account and learning from mistakes. Finally, we started managing portfolios on a larger scale using portfolio heuristics. We have never invested a single penny of our personal funds or client funds using an optimizer. We have always used heuristics and downside risk measures. The ultimate test of my confidence in portfolio heuristics is as follows: My mother-in-law's retirement portfolio is allocated using a risk-reward heuristic. I continue to have the same mother-in-law.

Summary and Conclusions

The problem of a "butterfly effect" when using a portfolio optimizer is very real. Simply using historic data with an optimizer is not an effective forecasting technique and therefore, not an effective portfolio management technique.

Techniques such as the ones covered in Michaud (1998) are effective. The Michaud book should be required reading for any practitioner interested in using a portfolio optimizer. This paper was concerned with broadening the range of possible solutions to the "butterfly effect" beyond Michaud (1998) to include the use of a simple risk/return heuristic to generate portfolios.

Finally, there are two principles covered in this paper that are important to the reader:

First, a simple model will always provide superior forecasts than a complex model.

Second, academics look backward and practitioners look forward. There will always be a conflict whenever a practitioner attempts to apply an academic explanatory model to a forecasting situation. Academic models are useful because of their discipline and insight into a system, but they are not necessarily meant to forecast.

Acknowledgements

The author would like to thank Thomas Connelly for his encouragement and comments.

Biography

David N. Nawrocki is professor of finance at Villanova University. He holds an MBA degree and a Ph.D. degree in finance from the Pennsylvania State University. Dr. Nawrocki is the author of a portfolio optimization package and is the director of research for The QInsight Group, an investment management firm. He can be reached through the following web sites:
http://www.handholders.com
http://www.homepage.villanova.edu/david.nawrocki
http://www.qinsight.com

References

Adler, Michael (1987). "Global Asset Allocation: Some Uneasy Questions." Investment Management Review, (September-October), 13-18.

Chopra, Vijay K. (1991). "Mean-Variance Revisited: Near-Optimal Portfolios and Sensitivity to Input Variations." Russell Research Commentary, (December).

Elton, Edwin J., Gruber, Martin J. and Manfred W. Padberg (1976). "Simple Criteria for Optimal Portfolio Selection." Journal of Finance, 31(5), 1341-57.

Elton, Edwin J., Gruber, Martin J. and Thomas J. Urich (1978). "Are Betas Best?" Journal of Finance, 33(5), 1375-84.

Elton, Edwin J. and Martin J. Gruber (1995). Modern Portfolio Theory and Investment Analysis, Fifth Edition, John Wiley and Sons, New York.

Frost, Peter A. and James E. Savarino (1988). "For Better Performance: Constrain Portfolio Weights." Journal of Portfolio Management, (Fall): 29-34.

Francis, Jack C. (1991). Investments: Analysis and Management, Fifth Edition, McGraw Hill, New York.

Hensel, Chris R. and Andrew L. Turner (1992). "Making Superior Asset Allocation Decisions: Implications of Recent Research Commentaries." Russell Research Commentary, (December).

Jobson, J. D. and Bob Korkie (1981). "Putting Markowitz Theory to Work." The Journal of Portfolio Management. (Summer): 70-74.

Kallberg, J.G. and William T. Ziemba (1984). "Mis-specification in Portfolio Selection Problems". In Risk and Capital, ed. G. Bamberg and A. Spremann, 74-87, Springer-Verlag, New York.

Markowitz, Harry (1959). Portfolio Selection: Efficient Diversification of Investments, John Wiley, New York. Second Edition, (1991), Basil Blackwell, Cambridge, MA.

Markowitz, Harry (1987). Mean-Variance Analysis in Portfolio Choice and Capital Markets. Basil Blackwell, Cambridge, MA.

Martin, A. D. (1955). "Mathematical Programming of Portfolio Selection." Management Science, 1(1), 152-66.

Michaud, Richard O. (1989). "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" Financial Analysts Journal, (January-February): 31-42.

Michaud, Richard O. (1998). Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation. Harvard Business School Press, Boston, MA.

Nawrocki, David N. (1983). "A Comparison of Risk Measures When Used in a Simple Portfolio Selection Heuristic." Journal of Business Finance and Accounting, 10(2), 183-94.

Nawrocki, David N. (1990). "Tailoring Asset Allocation to the Individual Investor." International Review of Economics and Business, 37(10-11), 977-88.

Nawrocki, David N. (1996). “Portfolio Analysis with a Large Universe of Assets.” Applied Economics, 28, 1191-98.

Nawrocki, David N. (1999). “A Brief History of Downside Risk Measures.” Journal of Investing, 8(3), 9-25.

Sharpe, William F. (1967). "A Linear Programming Algorithm for Mutual Fund Portfolio Selection." Management Science, (March): 499-510.