Portfolio Monitoring In Theory and Practice[1]

By

Richard O. Michaud, David N. Esch, Robert O. Michaud
New Frontier Advisors, LLC
Boston, MA 02110

Forthcoming in the *Journal Of Investment Management*

[1] We wish to acknowledge the helpful comments of the discussant Denis Chaves and other participants at the *Journal Of Investment Management* Fall Conference, Boston, 2010.

Abstract

The when-to-trade decision is a critical yet neglected component of modern asset management. Typical rebalancing rules are based on suboptimal heuristics. Rebalancing is necessarily a statistical similarity test between current and proposed optimal portfolios. Available tests ignore many real world portfolio management considerations. The first practical test for mean-variance optimality, the Michaud rebalancing rule, ignored the likelihood of information overlap in the construction of optimal and current portfolios. We describe two new algorithms that address overlapping data in the Michaud test and give examples. The method allows large-scale automatable non-calendar based portfolio monitoring and quadratic programming extensions beyond portfolio management.

Practitioner Summary

The when-to-trade decision is a critical yet neglected component of modern asset management. The need-to-trade decision is typically based on suboptimal heuristic rules. Until relatively recently no effective decision rule has existed for deciding whether a currently held portfolio has aged sufficiently to make trading desirable. The fundamental problem is that portfolio managers and financial theoreticians persist in ignoring the statistical nature of asset management and the impact of estimation error on effective decision making. The rebalancing decision is necessarily a statistical similarity test between the current drifted portfolio and a proposed new optimal. Many rebalancing rules in use have had little theoretical or statistical foundation and often lead to trading in noise or ineffectively using useful information. While statistical similarity tests are available in the financial literature none treat the real world portfolio management problem that requires inequality constraints, targeted risk portfolios, trading costs, and asset manager style customization. Inclusion of these necessary features requires compute-intensive resampling methods that overcome the limitations of familiar null distributions. The first practical similarity test for mean-variance optimality is the Michaud rebalancing rule. However, the original procedure ignored an often important consideration in that much of the information used to construct the current portfolio may be implicitly or explicitly included in the new optimal. This partial input match results in an overly conservative Michaud rebalance signal. We develop new algorithms that address overlapping data in the Michaud test. The new distribution defines a critical range for the Michaud rule and extends its applicability and power. We describe two procedures and give examples. The method allows large-scale automatable non-calendar based portfolio monitoring. The procedure is generalizable as a statistical similarity rule for quadratic programming contexts with potential applications well beyond portfolio management.

Keywords

asset allocation, mean-variance optimization, Markowitz optimization, Michaud optimization, Monte Carlo Simulation, Resampled Efficient Frontier, overlapping data, portfolio monitoring, quadratic programming, rebalancing, resampling, computational finance

Introduction

Modern asset management requires effective portfolio monitoring. The rebalancing decision is necessarily a statistical similarity test between the current drifted portfolio and a proposed new optimal. However, portfolio rebalancing in practice is typically based on suboptimal heuristics that often lead to trading in noise or ineffectively using useful information. Many managers routinely trade on a fixed calendar basis. Others estimate hurdle rates that ignore the portfolio context and statistical estimation issues (e.g., Masters 2003). Other approaches include trade optimality (Dybvig 2005), jumps in prices (Liu et al 2002), trading frequency (Leibowitz and Bova 2010), trading strategies (Perold and Sharpe 1989), and dynamic programming (Sun et al 2006, Markowitz and van Dijk 2003). None are based on portfolio statistical similarity.

Statistically based portfolio similarity procedures have been given by Shanken (1985), Jobson and Korkie (1985), and Jobson (1991) among others. However, the few statistical approaches require unrealistic investment assumptions in order to ensure familiar null distributions. Real-world portfolio management requires linear inequality constraints on portfolio weights, targeted risk levels usually on some variation of the Markowitz (1959) mean-variance (MV) efficient frontier, trading costs, and asset manager style customization.[2] Practical decision rules require compute-intensive resampling methods to create optimality tests which would be intractable using traditional analytical techniques.

The first practical procedure for statistical similarity is the Michaud rebalancing rule (Michaud 1998, extended in Michaud and Michaud 2008a, b). It measures the statistical similarity of the current portfolio relative to an associated MV risk targeted portfolio on the Resampled Efficient Frontier™ (REF).[3,4] In Michaud optimality, each portfolio on the REF is an average of properly associated resampled MV optimal portfolios. Michaud rebalancing compares the tracking error of a given portfolio to the target optimal relative to the distribution of tracking errors associated with the target optimal created by the

[2] Markowitz (2005) demonstrates the critical importance of realistic linear inequality as well as equality constraints on financial theory as well as applications.

[3] Michaud's REF is a generalization of the linear equality and inequality constrained Markowitz MV efficient frontier that includes estimation error in inputs and resampling to define portfolio optimality (Michaud and Michaud (2008a, Ch. 6, 2008b). The Michaud optimization and rebalancing test are protected by U.S. patents and patents pending.

[4] Michaud (and Markowitz) efficiency has been critiqued as inconsistent with the rationality axioms of the von Neumann and Morgenstern (1944) Expected Utility Hypothesis (EUH). However the EUH is itself subject to numerous serious critiques. As a consistent formal axiom system, EUH is neither consistent with observed investor behavior (Allais 1953, Kahneman and Tversky 1979), unique (Quiggin 1993, Luce 2000), nor complete (Gödel 1931, Church 1936, Turing 1936), and cannot claim definitive characterization of rational decision making under uncertainty. An alternative rational framework of investor behavior based on evolutionary principles (Farmer and Lo 1999, Lo 2004) can be viewed as fundamentally consistent with Michaud optimality. REF optimal portfolios properly defined enhance the likelihood of investor prosperity or survival in controlled statistical experiments relative to classical MV optimization (Michaud 1998, Ch. 6, Markowitz and Usmen 2003, Michaud and Michaud 2008a, Chs. 6, 9) in the investment period. In our case psychological hypotheses of species behavior or ecology are unnecessary. In practical terms, REF optimality is simply a convenient framework for constructing enhanced linear constrained MV efficient portfolios that address estimation error uncertainty endemic in applied finance and asset management.

resampling process under the statistical similarity null hypothesis.[5] A high percentile, or rebalance probability, indicates that the current portfolio is statistically far from target optimal and that trading is likely to be desirable. The threshold percentile level depends on the manager's risk aversion; i.e., whether or not the possibility of trading in noise is less desirable than ineffectively using investment information. The trading decision depends on the level of confidence associated with the manager's information, trading costs, and numerous additional investment and statistical issues.[6]

A critical consideration in portfolio monitoring is that much of the information used to construct the current portfolio may also be implicitly or explicitly reused when defining the new optimal. This partial input match results in an overly conservative Michaud rebalance signal. A portfolio similarity test that properly uses overlapping information corrects the statistical significance of the percentile measure while often recommending trading earlier or later than any fixed calendar timetable.

We develop new critical values that are conditional on a specified amount of common information in the Michaud test. These new thresholds define the appropriate critical ranges for the Michaud rule and boost its power. We describe two procedures, one for purely historical data and the other for the more common case of managed risk-return estimates.[7] The algorithms are illustrated with examples. We discuss customization of the monitoring rule in the context of manager, investor, and marketing imperatives. The method allows large-scale automatable non-calendar based monitoring. More generally the procedure provides a rigorous statistical context for quadratic programming estimation with potential applications beyond portfolio management.

The procedures we describe are addressed to the need-to-trade, not how-to-trade, decision. How-to-trade often requires consideration of many market, investment, and client factors and is beyond the scope of this report. We note that the trade-to portfolio may often be an intermediary point between current and target optimal.

A mathematical description of the portfolio monitoring framework is presented in Section I. A computational procedure adjusting Michaud's rebalancing test for overlapping data is given in Section II assuming optimization inputs are based purely on historical data. Later discussion for implementation in the more practical case of managed inputs is presented. Section III illustrates the procedures with applications to asset allocation. Section IV discusses considerations of using the rule in a practical monitoring scenario. Section V provides some generalizations, a summary and some conclusions.

I. Monitoring Framework

We begin by placing the REF, rebalance probability, and monitoring problem within a suitable mathematical framework. We index time periods starting with $t=0$ at the last

[5] Tracking error may be replaced by other measures of portfolio dissimilarity. For example, a utility-function-based metric could be used. We use tracking error as the most straightforward of these measures.
[6] We will discuss some of these issues further below.
[7] U.S. patent pending.

rebalance. Suppose information X_t is available at time t and is statistically modeled with density $f(X|\theta)$.[8] For the examples shown here θ represents the parameters $\{\mu, \Sigma\}$ of a multivariate normal model, although it need not be limited thus. The best available information for θ at time t comes from the posterior density $f(\theta|X_t)$, and the model for future observations X^* is the posterior predictive distribution with density $f(X^*|X_t)$. These two posterior distributions capture the manager's full set of information.[9] The REF optimal portfolio P_t at time t is the posterior expectation of a selected Markowitz Efficient portfolio[10] $P_{EF}(\theta)$ for known inputs θ. This is defined as

$$P_t = \int P_{EF}(\theta) f(\theta|X_t) d\theta.$$

Implicit in this definition is a method of selecting portfolios along the linear equality and inequality constrained efficient frontier for a given θ. At a particular value of the parameters $\{\mu, \Sigma\}$, portfolios can be chosen to maximize expected utility or by another method, but we refer to the set of portfolios chosen by any specific method as an associated set of MV optimal portfolios. Certain association schemes outperform others on simulated out-of-sample performance with regard to many typical and exemplary datasets taken from contemporary capital markets.[11] The best-performing association schemes should be preferred in practice. In practice the REF optimal portfolio is computed as the average of properly associated simulated MV efficient portfolios.

In order to quantify the rebalancing imperative, we must have a way to score suboptimal portfolios. Let $\mathcal{D}(P_1,P_2)$ be a discrepancy function between portfolios P_1 and P_2.[12] Then the predictive distribution $f(X^*|X_t)$ implies a probability distribution on $\mathrm{D}(P^*,P_t)$ where P^* is the REF evaluated on a random draw from $f(X^*|X_t)$. The Michaud rebalance probability $R(t)$ is defined as the cumulative distribution function of this implied distribution evaluated at the observed discrepancy $\mathrm{D}(P,P_t)$ between current drifted portfolio P and optimal P_t, i. e.

$$R(t) = \Pr(\mathcal{D}(P^*, P_t) < \mathcal{D}(P, P_t)).^{13}$$

[8] Return distributions are well known to be nonstationary. However, for many practical situations there is not enough data to support estimation of a more complex probability model, so the more parsimonious stationary model is preferred, and the input data is limited to a relevant time period to the analysis or weighted to decrease in importance for less relevant time periods. See Esch (2010) for further discussion.

[9] We use the traditional notations for Bayesian analysis but intentionally leave model choice at the discretion of the manager. Many types of information can be included in the analysis.

[10] Managers typically use linear inequality constraints when defining optimality in any time period. This practice is one of the main reasons analytical calculation of the rebalance test is intractable, and Monte Carlo methods must be used.

[11] See Michaud and Michaud (2008a, Ch. 6) and Esch(2012).

[12] The Michaud procedures by default use tracking error, i, e. relative variance, as the discrepancy function, but any suitable distance metric can be substituted. Any well-defined norm on portfolio space induces a discrepancy function when evaluated on the difference of portfolios.

[13] The Michaud procedures use the empirical form of the distribution of $\mathcal{D}(P^*,P_t)$ obtained by simulation. This is an approximation of the true continuous distribution which has no closed form for general inequality-constrained optimization problems.

An important practical limitation on $R(t)$, the central problem which this paper addresses, is that much of the information in X_0 is also part of X_t. This information overlap creates a downward bias in the calculation of $\mathcal{D}(P,P_t)$ when it is measured against independent draws of $D(P^*,P_t)$ based on the predictive distribution. The method presented in Section 2 will overcome this limitation by including the overlapping information in the simulation process.

To mathematically define the central problem and solution of this paper we must first introduce parameters L, k, and T. The parameter L defines the nominal critical value of the rebalance test, or the percentage of the distribution of $\mathcal{D}(P^*,P_t)$ which the manager wishes not to exceed. The parameter k defines the normal rebalancing period for a given manager. For example, a value manager may anticipate the need to rebalance every twelve months while more active managers may anticipate rebalancing more frequently. The parameter T represents the number of total independent observations in the information set used to define optimality. Portfolios are optimized from many information sources and only expected return and covariance matrices are typically available, so both T and k can be heuristically adjusted to match the investor's preferences.

The decision to trade is often a function of accumulated new data. Usually the optimization to calculate P_t at time $t > 0$ uses much of the same information used to construct P. Due to overlapping data, rebalance probabilities may appear quite small, yet appropriately signal a need to trade in spite of being substantially less than the nominal threshold L.

Our overlap-adjusted threshold for the Michaud rebalance test is denoted as $C_L(k)$. We define $C_L(k)$ as follows: let X_t be composed of components $X_{t\cap 0}$ and $X_{t\backslash 0}$, the intersection and set difference of X_t with X_0, respectively. Then $C_L(k)$ is defined as the L^{th} percentile of the distribution of $R(t)$ induced by replacing $X_{t\backslash 0}$ with a random draw from its predictive distribution within the calculation of P_t.

The next section describes general computational procedures for calculating $C_L(k)$. $C_L(k)$ has limit L for increasing k when P and P_t are associated portfolios on the REF. As t increases, $R(t)$ increases and crosses the critical probability $C_L(k)$ at some time t. In the context of overlapping data in the optimization inputs, $C_L(k)$ is the lower limit of a more suitable range for interpreting a rebalance statistic as a positive rebalance signal at confidence level L. $C_L(k)$ provides a rigorous statistical benchmark for a given data set and parameters L and k for determining whether trading is desirable at a given time period.

II. Computing $C_L(k)$

We address two cases for calculating $C_L(k)$. The first, described in detail, assumes that the MV inputs associated with both P and P_t and for computing REF portfolio optimality are defined entirely by prior historical returns. This case is not very practical but useful in clarifying the main features of the monitoring computation. The second case generalizes the framework for many investment processes in practice.

A. MV optimization inputs from T periods of historical returns.

Return series $X_0 = [x_1, \ldots, x_T]$ is used to calculate optimization inputs associated with P. A sequence of simulated optimization inputs are created for time $t=k$, also based on T historical time periods. The following algorithm uses a partial resampling framework to simulate overlapping datasets and obtain the empirical distribution of $R(t)$. $C_L(k)$ is computed for the given data set and parameters k, T, L, and Z, where Z is the number of Monte Carlo simulations of rebalance probabilities. The definition of a REF optimal portfolio at time t also requires an assumption, called the Forecast Confidence™ (FC) parameter that defines the number of simulated observations used in computing simulated mean-variance efficient frontiers in the portfolio averaging process. While the default assumption in this report is to assume $FC=T$, we discuss alternatives and applications later in the text.[14]

Algorithm A: Computing $C_L(k)$ from known historical returns[15]

- Compute P_0, the optimal target portfolio at time 0, from X_0.
- For i in $\{1, \ldots, Z\}$:
    - Replace k randomly selected returns[16] in X_0 with k returns simulated from the predictive distribution $f(X^*|X_t)$ and compute the optimal portfolio P_k^i
    - Compute the rebalance probability $R^i(k)$ of P_0 with respect to P_k^i
- $C_L(k)$ is L^{th} quantile of $R^i(k)$

B. General MV optimization input estimation.

Few, if any, asset managers rely solely on historical returns for computing MV estimates for actual investment. Risk-return estimates generally reflect information from a wide variety of sources and aspects of the investment process.[17] REF optimized portfolios P_0 and P_t are based directly on inputs calculated at times 0 and t. In this context there is no explicit value of T to define the number of draws from the input distribution; consequently its value must be assumed. For this case, T represents the manager's

---

[14] The patented FC procedure can be conveniently formulated as an index with values from one to ten to heuristically reflect increasing certainty in risk-return point estimates independent of T. FC can be associated with numerous practical investment management issues. See further discussion in Michaud and Michaud (2008a,b).

[15] Chaves (2010) suggests a computationally efficient approximation to our algorithm. This approach modifies the calculation of $R(t)$ but has some important limitations. In our method, each scenario has its own updated posterior and discrepancy function. Simulating separate rebalance tests provides a more realistic calibration to overlapping data, as well as providing additional flexibility with settings of various parameters of the optimization process, as detailed in Section 5. Managers may also need flexibility to use different predictive distributions conditioned on different information sets in the simulation process from the rebalance test so as not to reflect unusual events in the contrafactual simulations.

[16] See Davison and Hinkley (1997) for discussions and generalizations on simulation schemes.

[17] In practice, asset managers use many techniques for modifying MV inputs and enhancing their forecast value. See Michaud and Michaud (2008a, Chs. 8 and 11) for further discussion and references.

assumption of the information cycle or information level associated with the number of draws of independent, identically distributed (iid) returns at time t. The adjusted critical values $C_L(k)$ in this algorithm (B) are computed by applying Algorithm A to simulated datasets in place of the observed data X_0.[18] Since some information is lost in reducing the dataset to optimization inputs, the two analyses may produce somewhat different results. We have indeed observed variability in the results from the above algorithm on multiple datasets simulated from a single set of inputs, as shown in Figure 1.

III. Examples

We provide a comprehensive illustration of the new procedure with respect to a data set used in prior REF optimization studies. This data set represents an asset allocation case for eight capital market indices as described in Michaud (1998, Ch. 2)[19]. The main results are shown as distribution plots and percentile curves in Figures 1 and 2, and detailed results of a few modifications to the inputs appear in the appendix.

The Michaud (1998, Ch.2) data consists of 18 years of monthly historical total returns for eight capital market indices, six equities and 2 fixed income, from January 1978 to December 1995. For our main analysis we use T=120 monthly periods; we also treat T=60 in the appendix. Unless indicated we also assume $FC = T$. P_t for this case would be the REF optimized portfolio estimated at t monthly periods subsequent to December 1990. For example, if $t = 12$, then P_t is the REF optimized portfolio at December 1991 calculated using the previous 120 months of data. Adjusted critical values $C_L(k)$ are computed for values of k from 1 to 24.

Figures 1 and 2 display two different presentations of the critical rebalance probability as a function of k for the middle portfolio on the REF based on ten years of monthly historical returns for the indicated confidence levels.[20] The middle of the box plots in Figure 1 represent the 50% value of L and the extremes the 25% and 75% level, for $C_L(k)$ for each value of k for the indicated cases. The display also shows that the need-to-trade probability required with a one-year rebalance cycle (k = 12) at the 90% confidence level for T=120 is roughly 20%. Figure 2 simply traces $C_L(k)$ as a function of k for different levels of L. Theoretically $C_L(k)$ is monotone increasing as a function of k; Monte Carlo

---

[18] The computational algorithm and further technical details are given in Michaud et al (2010).

[19] We used the same data set that has been used in key articles written on Michaud optimization. These include Michaud (1998), Michaud and Michaud (2008a), Markowitz and Usmen (2003), Harvey, Leichty, and Leichty (2008), and Michaud and Michaud (2008c). This data set is sufficiently complex to illustrate our new methods. Complete tables of means, standard deviations, correlations, and many efficient frontier portfolios are available in the above references. Our qualitative results hold for a wide range of datasets.

[20] Due to compute-efficiency and statistical stability considerations, the return-rank algorithm is the recommended procedure in Michaud (1998), Michaud and Michaud (2008a,b) for computing REF optimal portfolios. Return ranks identify increasing risk levels of portfolios on the REF. As in Michaud and Michaud (2008a,b), we compute 51 return-rank portfolios, from low to high, on the REF. We identify portfolio 26 as the middle portfolio on the REF. There is no loss in generality if quadratic utility, stock/bond ratios, tracking error relative to a benchmark, or other procedures are used to identify different risk levels on the REF. As we will show, we have found that the statistical character of $C_L(k)$ may not vary much for different risk levels across the frontier.

estimation error is responsible for deviations.[21] In effect, $C_L(k)$ is simply a recalibration of the need-to-trade probability for given L.

The simulated rebalance probabilities shown in Figures 1 and 2 have several noteworthy properties. Both figures vividly demonstrate the impact of overlapping data on the Michaud rebalancing rule. Even after a year, the critical probability required for trading to optimal is far less than confidence level L. Figure 1 also shows that the critical probability, or benchmark for measuring optimality relative to estimation error and assumptions, has less variation in Algorithm A than B due to more specificity in starting assumptions.

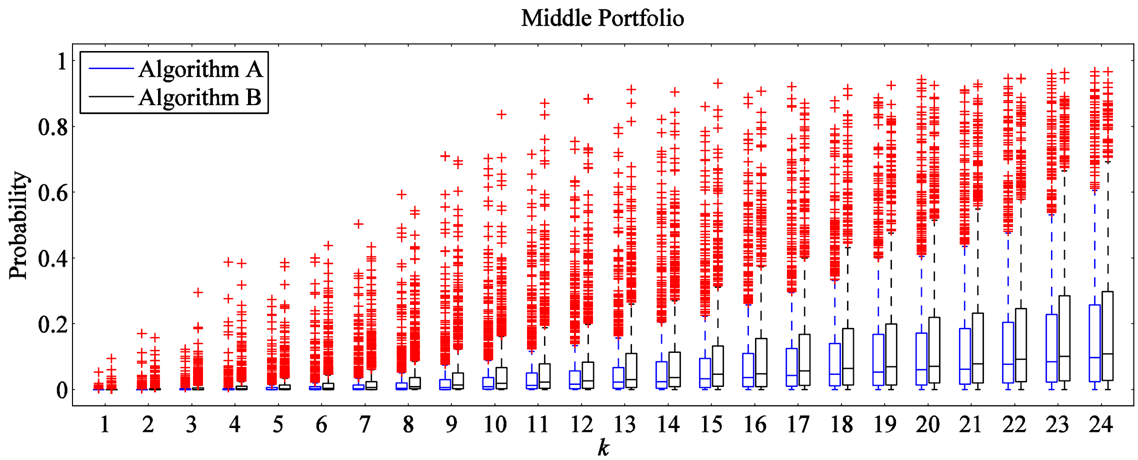

Figure 1: Michaud (1998) data. Michaud rebalance probabilities are simulated, replacing k months of 120 in each simulation, with T_0 = December 1990. Each combination of parameters was run with 1000 simulations. Boxplots (Tukey 1977) show distribution summaries of simulated Michaud rebalance probabilities for values of k from 1 to 24. Algorithms A and B are juxtaposed for each value of k. The three horizontal rules in the boxes mark the three quartiles of the distribution, and the dotted line whiskers extend to 1.5 times the interquartile range, or the limit of the data if it is closer. Values beyond the whiskers are marked with "+" symbols.

The appendix provides further insights into the statistical character of the adjusted critical values $C_L(k)$. Rather surprisingly, there is little variation across the different risk levels of the efficient frontier. Although the rebalance tests are comparing greater tracking errors for greater target risk levels on the frontier, the relative positions of the simulated aged portfolios among the simulated statistically equivalent portfolios remains remarkably consistent for otherwise equivalent cases. Variation of the number of simulated periods of overlap, via the parameter T changes the rate at which the $C_L(k)$ converges to its limiting value as k increases. The overlapping data effect vanishes more quickly for increasing k as T increases.

Other changes in input can affect the limiting distribution to which the simulations converge. Changing target rank between times 0 and t or using a number of FC periods different from T will affect the value of $C_L(k)$. Details of these effects appear in the appendix.

[21] Calculating $C_L(k)$ for many values of k is particularly computationally intensive. In applications managers may be interested in $C_L(k)$ at only one value of k. For our purposes, the computation provides a comprehensive illustration of the procedure.

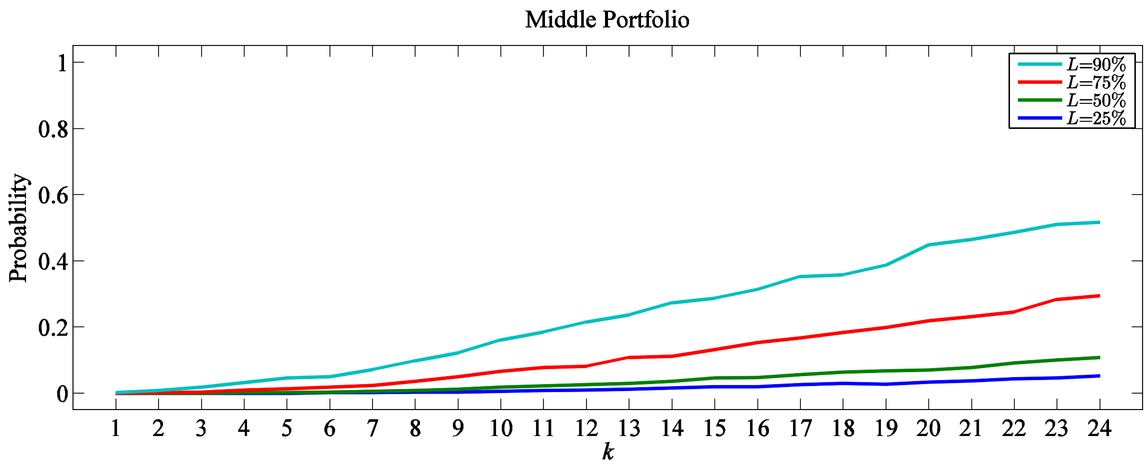

Figure 2: $C_L(k)$ as a function of k from the Algorithm B data in Figure 1. Here we show the 25th, 50th, 75th, and 90th percentiles of the distributions.

IV. Implementing the Michaud monitoring rule

Implementation of the monitoring rule requires appropriate settings of the parameters. Proper tuning of L, FC, k, and T allows customization of the procedure relative to a manager's investment process, client objectives, outlook, and other considerations.

The parameter L quantifies trading risk aversion. A large value of L reduces trading frequency by avoiding the danger of trading in noise but may also ignore performance opportunities. Alternatively, a small L increases trading frequency by more closely tracking to target but may result in statistically ineffective trades. A default value of L=.5 may often be useful.

The appropriate value of L also depends on a manager's assessment of the level of information in risk-return estimates at given time t, which may vary depending on various investment issues. These considerations are also reflected in the choice of the forward looking FC parameter which is embedded within the posterior and posterior predictive distributions $f(\theta|X_t)$ and $f(X|X_t)$. Higher FC levels create less dispersion in the resampled datasets, and generally more active use of investment information in the REF portfolios. The choice of FC may interact with L. For example, a high FC value may be associated with a lower level of L and vice versa. This is because the manager may decide that the information is very reliable and tracking to optimality is more important than the likelihood of trading ineffectively.

Setting the value of T is straightforward when the only source of information is historical data as in algorithm A. However, because risk-return estimates are typically managed, practical application will often require Algorithm B and an estimate of T. In this case T does not necessarily reflect a historical time period but more appropriately the denominator in the fraction k/T of new information in the analysis. Alternatively, T reflects the time it takes for information to completely cycle out of the analysis. As T becomes smaller, the rebalancing threshold $C_L(k)$ will increase more rapidly as k increases.

Once parameters have been set and $C_L(k)$ computed, the implementation of the monitoring rule is straightforward. The manager calculates the Michaud rebalance test $R(t)$ for the monitored portfolio P relative to the new REF optimal P_t and compares the value to $C_L(k)$. Rebalance probabilities greater than $C_L(k)$ indicate that rebalancing may be desirable. Note that $C_L(k)$ is simply a new scale for L. It may often be convenient to recalibrate the $C_L(k)$ distribution onto the L scale. The parameters defining the monitoring rule provide interesting opportunities to dynamically manage the trading decision in real time.[22]

In applications the set of monitored portfolios may differ by risk level. The Michaud monitoring rule can be accommodated by associating each portfolio with one of several pre-defined target REF portfolios. Once the critical probabilities are computed, the monitoring rule requires only a table lookup. Consequently automated customizable portfolio monitoring may be practical even for very large-scale applications.

V. Summary and Extensions

The when-to-trade decision is a critical function of professional asset management. Trading on noise is costly and unproductive but trading too little uses investment information suboptimally. Trading effectiveness impacts investment performance and manager competitiveness. At a macro level, trading effectiveness affects efficient allocation of capital in global markets.

Yet current practice is largely characterized by suboptimal calendar rules and simple point estimate heuristics. Portfolio monitoring is essentially a statistical test of the similarity of two portfolios. Until recently, statistical similarity tests required unrealistic assumptions. The Michaud rebalancing test addressed practical portfolio construction and investment management considerations but is overly conservative if the two portfolios are based on overlapping data. We describe two new algorithms for properly determining when trading is statistically desirable with overlapping information.

The procedures we present have a number of generalizations. In particular there is no need for the target portfolio to reside on the REF or any MV efficient frontier. Since the ability to resample datasets enables computation of statistical equivalence regions, any portfolio calculation method, along with a measure of portfolio similarity, can produce a rebalance test statistic which could be calibrated by the procedure in this paper. Note also that θ can parameterize any return distribution, not just mean-variance based.

While our methods and illustrations are inspired and guided by the demand characteristics of modern asset management in practice, they are not limited to this context. The procedures are applicable to a wide variety of process monitoring applications where inequality or other constraints require computational methods for estimation and inputs include overlapping data. In particular the procedures can be used for data analysis monitoring applications that uses multivariate linear regression with overlapping data and

[22] For example, the VIX index may indicate changes in market volatility that may impact need-to-trade decisions.

linear inequality constraints. More generally, the procedure may be viewed as providing a statistical context for many quadratic programming applications.[23]

Our portfolio monitoring framework presents novel concepts and challenges for managers and investors. However, the effort to quantify the monitoring decision with properly defined parameters may be helpful in clarifying and controlling investment process and intuition. Investors may find such a manager likely to have enhanced performance over traditional ones. The potential for adding investment value and enhancing reliability seems a challenge worth the effort.

[23] Theil (1971, pp. 347-54) notes that analysts have or should have a view of the range of values of coefficients in multivariate linear regression estimation that should be included in the procedure. He also notes that quadratic programming may have important applications in proper data analysis if a statistical basis for the procedure were available.

Appendix
Detailed results for various modifications to the inputs

Consistency across the Frontier

Figure 3 displays the simulations of Figure 1 for low to high risk on the entire efficient frontier broken out by single values of k on each panel.[24] Although there is a slight increase in the probabilities needed for trading at the lower and upper ends of the efficient frontier, the distribution of simulated probabilities is remarkably consistent across different risk levels of the frontier. Most if not all of our experimentation has shown the same result: the target risk level on the frontier does not greatly affect the distribution of simulated rebalance probabilities, and hence $C_L(k)$, despite the wide variation in the distribution of tracking errors used to calculate each rebalance test.

Number of Input Time Periods

Figure 4 displays an analysis identical to Figure 1 except that it is based on T=60 monthly historical returns rather than 120. Comparison of the two figures demonstrates that an increase in estimation error in risk-return estimates significantly increases the critical probability required for trading to optimal. This increase occurs because each additional simulated return as k is incremented has twice the information relative to the whole data set. Note that even in this case, the critical probability is small relative to confidence level at short monitoring periods reflecting the dominant effect of overlapping data. On the other hand, the overlapping data effect tends to diminish rapidly much beyond a year and at k=24 the distribution of rebalance probabilities is much closer to uniform than in the T=120 case.

[24] As discussed earlier, the results are computed from low to high risk computed from the return-rank algorithm for 51 points on the frontier.

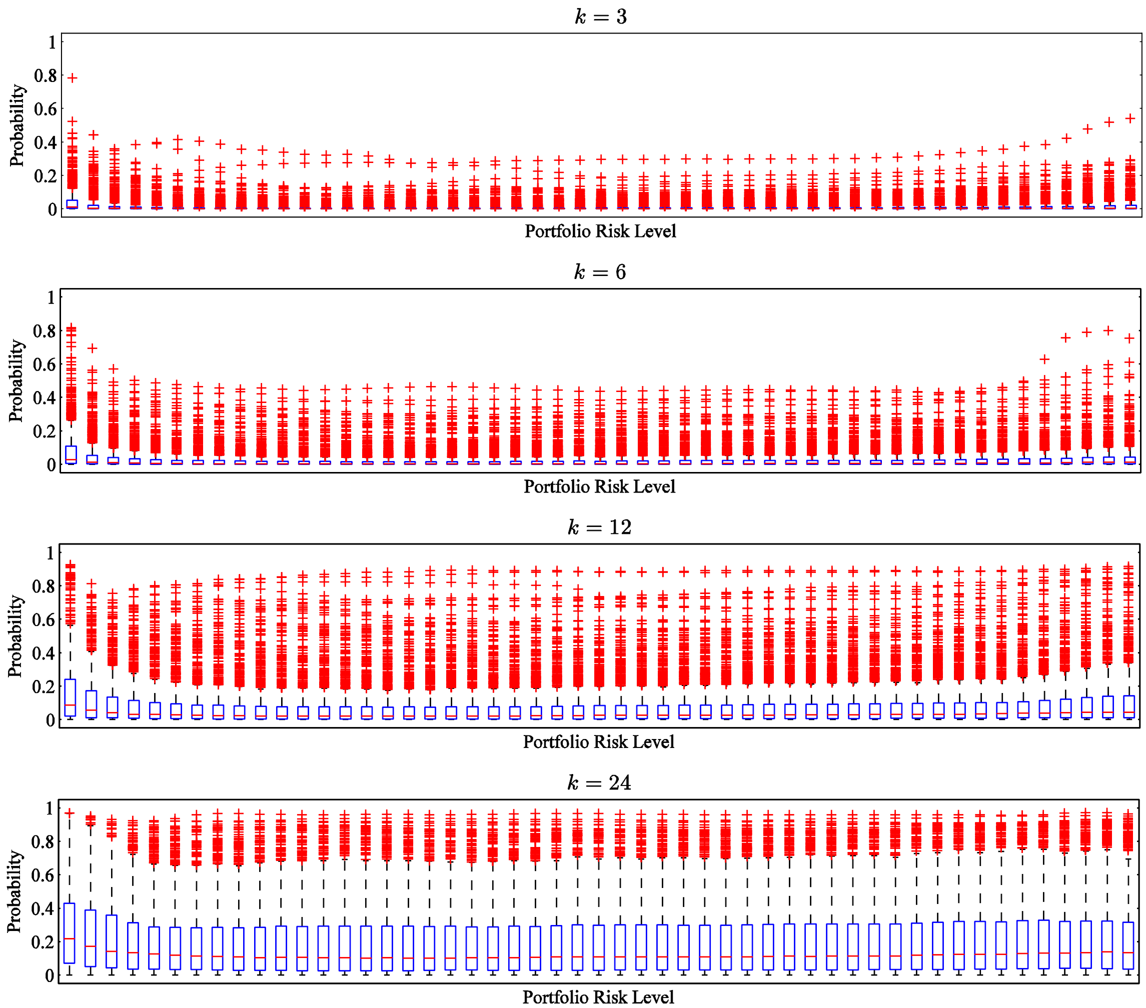

Figure 3: The Algorithm B simulations are shown for given ***k*** across the entire range of frontier risk levels in each panel. This Chart shows surprisingly little variation across different risk levels in the rebalance statistic, except at the ends of the frontier. There is variation across the frontier within resampled datasets, as can be seen by the paths of some outliers, but the aggregated simulations occupy a remarkably stable range across the frontier.

Non-associated portfolios at t=0 and t=k

Risk preferences may change over time. In such cases target portfolio risk may vary from one period to another, and the analyst would choose a different set of associated portfolios to average when computing the REF. Changes in portfolio risk alone are often enough to indicate that rebalancing may be desirable. When a different frontier point is targeted, the analyst should use the value of $C_L(k)$ computed such that P_0 and P_k are associated with P_t. This applies a uniformly strict standard to the rebalance test regardless of previous status.

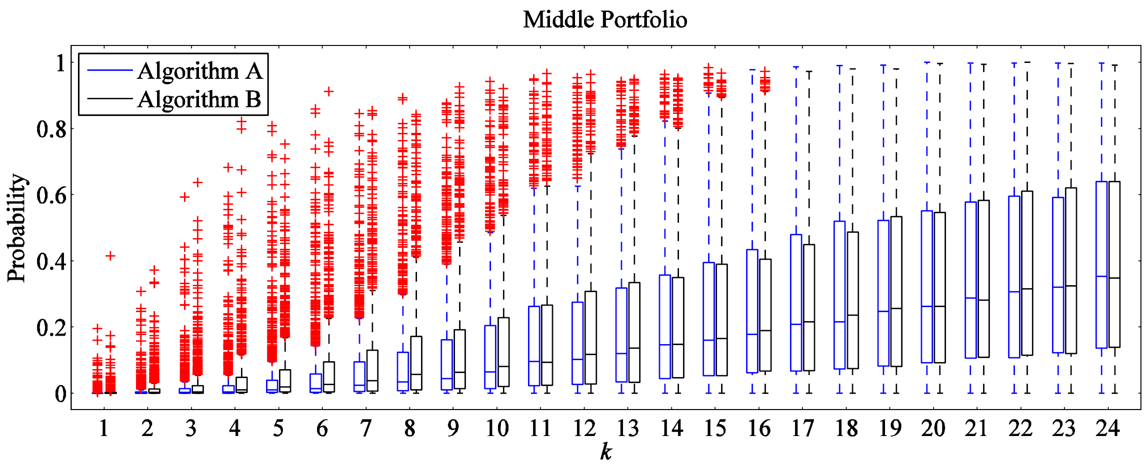

Figure 4: Boxplots for Michaud (1998) data, Algorithms A and B, T=60. Note the much more rapid convergence of the distributions towards uniform coverage of the [0,1] interval.

While not useful as strict criteria for rebalancing, the distributions of overlap-simulated rebalance probabilities from different parts of the efficient frontier are useful in understanding the behavior of the rebalance test and how changing target risk affects the need to rebalance. The upper panel of Figure 5 illustrates the distributions of rebalance probabilities for lower target risk P_k relative to a middle REF optimized portfolio P_0.[25] For comparison, the lower panel of Figure 5 shows the identical analysis when both P_0 and P_k are middle target risk, and is identical to the Algorithm B data of Figure 1. In this figure the distribution of rebalance probabilities is shifted towards greater values in the upper panel. When target portfolio risks are different the L^{th} quantile of the distribution of overlap-simulated rebalance probabilities in fact tends toward a value greater than L as k approaches T. Only a small probability mass of the distributions in the upper panel of Figure 5 occur below the medians in the lower panel. This fact is consistent with the intuition that changing portfolio target risk alone often requires rebalancing.

Variations in *FC*

In prior examples it was convenient to assume that the T and FC periods are the same. However, in applications, each parameter has a separate role to play in defining the monitoring process. Their roles are clearest when considering historical data as in algorithm A. In this case T is the number of periods in the data set while FC is the number of resampling periods representing the investor's level of certainty in investment information. While T is fixed FC may vary depending on market outlook, manager style, investment horizon, and other considerations.

When T and FC differ, $C_L(k)$ no longer necessarily converges to L as k increases. This is because the two parameters imply different amounts of information and levels of dispersion about the target portfolio. The net effect is that $C_L(k)$ tends toward values smaller or larger than L depending on whether the number of FC resampling periods is

[25] The lower risk target optimal portfolio relative to the middle portfolio is computed for rank 16 of 51 portfolios with equally-spaced expected returns.

less or greater than T, respectively. Figure 6 shows the effect of changing the FC window from 120 to 60 time periods with T fixed at 120 time periods. Since 120 time periods yield more information and less dispersion of simulated portfolios than 60 periods, the simulated rebalance test statistics are small and $C_L(k)$ converges to relatively small values.

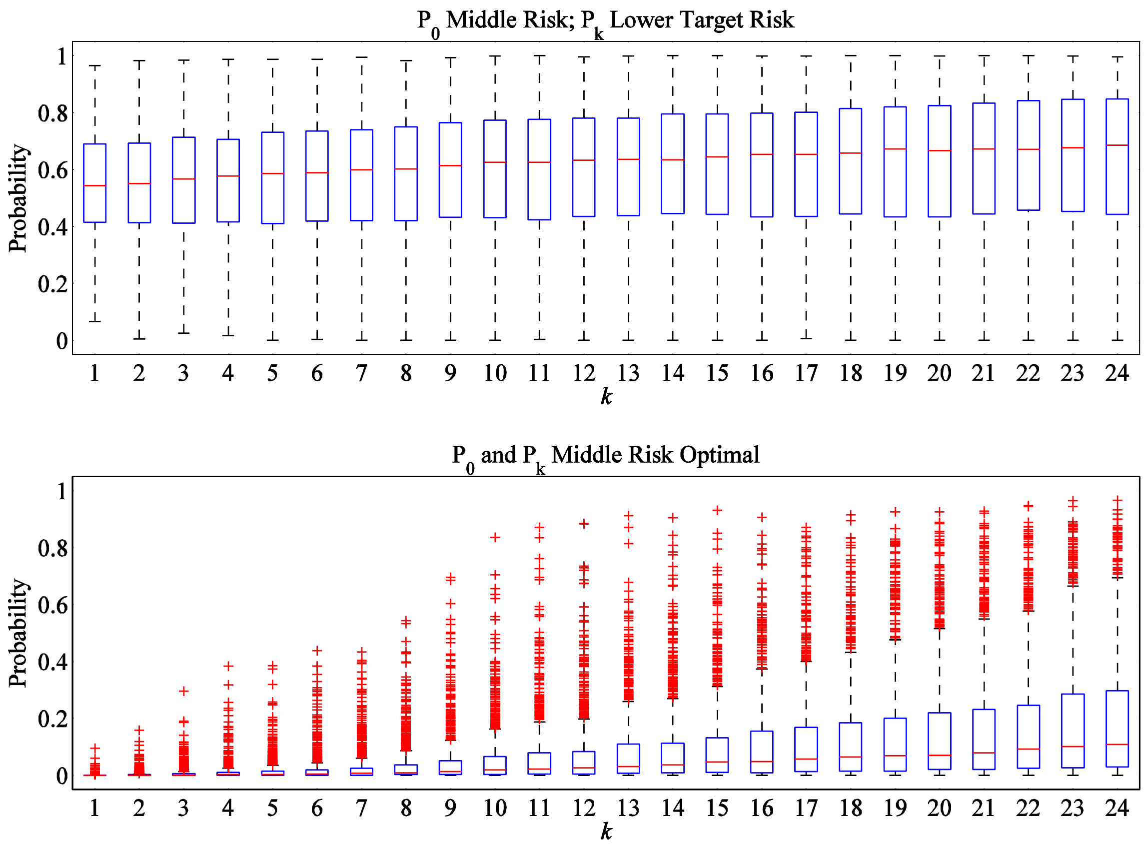

Figure 5: Algorithm B, lower target risk level. The data in panel 2 exactly matches the algorithm B data from Figure 1. When the target risk differs from initial, the simulated rebalance statistics are typically greater than when they match, and converge to a nonuniform distribution as $\boldsymbol{k}$ increases. The boxplots for large $\boldsymbol{k}$ in the first panel can be seen to have quartiles at values greater than 25%, 50%, and 75%, the quartiles of the uniform distribution.

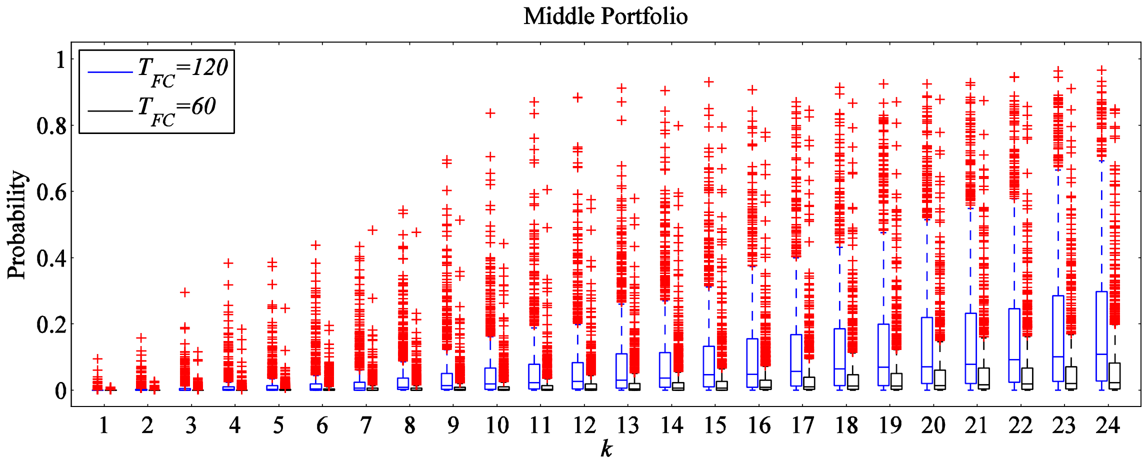

Figure 6: An illustration of the effect of decoupling the Forecast Confidence™ from the total information parameter T used in computing $C_L(k)$. The data are taken from Michaud (1998), using 120 months of data. The left series of boxplots show the same (Algorithm B) data of in Figure 1, for which both **FC** and total information T were set to 120 time periods. The right series shows the effect of changing the FC to 60 time periods. In this case $C_L(k)$ is converging to values less than L since replacing a set of 120 observations results in less dispersion about the optimal portfolio than does replacing only 60 observations.

References

Allais, M. 1953. "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine." *Econometrica* 21: 503-546.

Chaves, D. 2010. Discussion of Michaud et al. "Dynamic Portfolio Monitoring." *Journal Of Investment Management* Fall Conference, Boston.

Church, A. 1936. "An unsolvable problem in elementary number theory." *American Journal of Mathematics* 58:345-363.

Davison, A.C., and Hinkley, D. V. 1997. *Bootstrap Methods and their Applications*. New York: Cambridge University Press.

Dybvig, P. H. 2005. "Mean-variance portfolio rebalancing with transaction costs." Working paper, Washington University in Saint Louis.

Esch, D. 2010. "Non-Normality Facts and Fallacies." *Journal of Investment Management* 8 (1):.49-61.

Esch, D. 2012 "Averaging Frontiers." Working Paper, New Frontier Advisors, LLC.

Farmer, D. and A. Lo 1999. "Frontiers of Finance: Evolution and Efficient Markets." *Proceedings of the National Academy of Sciences* 96, 9991-9992.

Gödel, K. 1931. "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme." (On Formally Undecidable Propositions in Principia Mathematica and Related Systems). *Monatshefte für Mathematik und Physik* 38:173-98.

Harvey, C. Liechty, J. and Liechty, M. (2008). "Bayes vs. Resampling: a Rematch." *Journal Of Investment Management* 6 (1): 1-17.

Jobson, D. and Korkie, B. 1985. "Some Tests of Linear Asset Pricing with Multivariate Normality." *Canadian Journal of Administrative Sciences* 2: 116 - 140.

Jobson, D. 1991. "Confidence Regions for the Mean-Variance Set: An Alternative Approach to Estimation Risk." *Review of Quantitative Finance and Accounting* 1:235-257.

Kahneman, D. and Tversky, A. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47: 263-291.

Leibowitz, M and Bova, A. 2010. "Slow Rebalancing." *Morgan Stanley Research*, July.

Liu, J., Longstaff F., and Pan, J. 2003 "Dynamic Asset Allocation with Event Risk." *Journal of Finance* 58 (1): 231-259.

Lo, A. 2004. "The Adaptive Markets Hypothesis." *Journal of Portfolio Management* 30:15-29.

Luce, R.D. 2000. *Utility of Gains and Losses*. Mahwah, NJ: Lawrence Erlbaum Associates.

Markowitz, H. 1959. *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley. 2nd ed. Cambridge, MA: Basil Blackwell, 1991.

Markowitz, H. and Usmen, N. 2003. "Resampled Frontiers vs Diffuse Bayes: An Experiment." *Journal Of Investment Management* 1(4).

Markowitz, H., and Van Dijk, E. L. 2003. "Single-period mean-variance analysis in a changing world." *Financial Analysts Journal* 59 (2): 30-44.

Markowitz, H. 2005. "Market Efficiency: A Theoretical Distinction and So What?" *Financial Analysts Journal* 61(5):17-30.

Masters, S. J. 2003. "Rebalancing." *Journal of Portfolio Management* 29 (3): 52-57.

Michaud, R. 1998. *Efficient Asset Management*. New York: Harvard Business School Press. Now published by Oxford University Press.

Michaud, R. and Michaud, R. 2008a. *Efficient Asset Management*. 2nd Ed. Oxford University Press: New York.

Michaud, R. and Michaud, R. 2008b. "Estimation Error and Portfolio Optimization: A Resampling Approach." *Journal of Investment Management* 6 (1):8-28.

Michaud, R. and Michaud,, R. 2008c. "Discussion on article by Campbell R. Harvey, John C. Liechty and Merill W. Liechty, Bayes vs. Resampling: A Rematch (2008, Volume 6, Number 1)." *Journal of Investment Management* 6 (3):1-2.

Michaud, R., Esch, D. and Michaud, R. 2010. "Technical Note on Dynamic Portfolio Monitoring." New Frontier Advisors, September.

Perold, A. and Sharpe, W. 1988. "Dynamic Strategies for Asset Allocation." *Financial Analysts Journal* 44 (1): 16-26.

Quiggin, J. 1993. *Generalized Expected Utility Theory*. Boston: Kluwer.

Shanken, J. 1985. "Multivariate Tests of the Zero-Beta CAPM." *Journal of Financial Economics* 14(3): 327-357.

Sun, W., Fan, A., Chen, L., Schouwenaars, T., Albota, M. 2006. "Optimal Rebalancing for Institutional Portfolios." *Journal of Portfolio Management* 32 (2): 33-43.

Theil, H. 1971. *Principles of Econometrics*. Wiley: New York.

Tukey, J. W. 1977. *Exploratory Data Analysis*. Addison-Wesley.

Turing, A.M. 1936. "On Computable Numbers, with an Application to the Entscheidungsproblem." *Proceedings of the London Mathematical Society* 2 (42): 230–65.

von Neumann, J. and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, N.J.