Journal Of Investment Management, Vol. 1, No. 4, (2003), pp. 9–25

JOIM
www.joim.com

RESAMPLED FRONTIERS VERSUS DIFFUSE BAYES: AN EXPERIMENT

Harry M. Markowitz[a,*] *and Nilufer Usmen*[b]

Introduction

This paper reports on an experiment which compares two methods of handling the fact that empirically observed means, variances, and covariances, for a mean–variance analysis, are themselves noisy. One method is Bayes inference using diffuse priors which the present authors, among many others, have recommended. (Markowitz and Usmen, 1996a,b). The other is the method of Resampled Efficient Frontiers™ recommended by Richard O. Michaud (Michaud, 1998).[1]

The experiment is a computer simulation "game" with two players and a referee. In the game the referee generates 10 "truths" about eight asset classes. For each truth the referee draws 100 different possible "histories" of 216 monthly observations. (We chose eight asset classes and 216 months to keep the experiment as close as possible to that of Michaud.)

Each history is presented to each player. The players know that the truth is a joint normal distribution with unchanging means, variances, and covariances but do not know the parameter values. The Michaud player uses the observed history to generate a resampled frontier. That is, for a given history the player randomly generates many mean–variance efficient frontiers and averages these. The Bayes player uses the observed history to update beliefs, from prior to posterior, then uses these beliefs to compute one efficient frontier. Because of the high dimensionality of the "hypothesis-space," Monte Carlo sampling must be used to approximate the Bayes player's *ex post* means, variances, and covariances. Given their respective frontiers each player picks three portfolios, namely, the portfolios which each player believes maximizes

$$EU = E - \lambda V \tag{1}$$

for $\lambda = 0.5, 1.0, 2.0$, where E and V are the portfolio mean and variance. The referee notes the player's actual expected utility using the true means, variances, and covariances—known only to the referee. The referee also notes each player's estimate of its expected utility. This is repeated for the 100 randomly drawn histories for a given truth and the 10 truths of the game.

[a]Harry Markowitz Company, San Diego, California, USA.
[b]School of Business, Montclair State University, Upper Montclair, New Jersey, USA.
*Corresponding author. Harry Markowitz Company, 1010 Turquoise Street Suite 245, San Diego, CA 92109, USA. Tel.: (858) 488-7212.

The assumption of normality and unchanging distributions may be unrealistic, but both players are apprised of the rules of the game. It is not obvious that the assumptions favor one methodology over the other. The authors expected the Bayesian approach with diffuse priors to do better than the resampled frontier approach. In fact, the opposite turned out to be the case. Section 1 of this paper describes how the referee generates truths and, from these, the histories "observed" by the players; Section 2 describes the actions of the Michaud player; Section 3 describes the actions of the diffuse Bayesian player; Section 4 presents the results of the experiment; Section 5 points out some questions raised by these results; Section 6 summarizes.

1 The referee and the game

The experiment ("game") is outlined in Exhibit A. The referee generates 100 histories from 10 "truths," each history consisting of returns on eight asset classes during 216 consecutive months. Each truth is itself randomly generated by the referee by computing the means, variances and covariances of 216 draws of eight returns each from a "seed" distribution. This seed distribution is normally distributed with means, variances, and covariances equal to the historic excess return over the US 30-day T-bill rate of the eight asset classes listed in Table 1 for the 216 months from January 1978 through December 1995, as in Michaud (1998).

Exhibit A The Experiment.

Referee chooses First/Next "Truth"

"Truth" is a joint normal return distribution with fixed mean vector μ and covariance matrix C not known to the players.

Referee draws First/Next historical sample randomly from Truth

For Player = {Bayesian, Resampler}

Referee gives historical sample to Player.

Player applies its procedure to sample. (See write-ups of respective procedures.)

For the given sample and for each utility function (specifically, for $EU = E - \lambda V$ for $\lambda = \frac{1}{2}$, 1, and 2) the Player returns:

Selected Portfolio
Estimate of its Expected Utility

For each (Player, Utility function):

Referee computes True expected utility.

Repeat for Next Historical Sample

After all historical samples have been generated and processed, and with Truth still fixed:

For each utility function, see which player had higher *EU* on average.
Compare *EU* achieved versus *EU* anticipated on average.

Repeat for Next Truth

Did one of the players do better for most Truths or on average?

Table 1 Asset classes used in experiment.[a]

Asset class	Data source
Canadian Equities	Morgan Stanley Capital International[b]
French Equities	Morgan Stanley Capital International[b]
German Equities	Morgan Stanley Capital International[b]
Japanese Equities	Morgan Stanley Capital International[b]
United Kingdom Equities	Morgan Stanley Capital International[b]
United States Equities	S & P 500 Index total return
United States Bonds	Lehman Brothers[c]
Euro Bonds	Lehman Brothers[d]

[a]Source: Michaud (1998) p. 13, footnote 16.
[b]Dollar return indexes net of withholding taxes.
[c]Government/Corporate US bond index.
[d]Eurobond global index.

Having thus established a truth, the referee generates a 216 month "history" from this truth by sampling joint normally from the truth's mean vector and covariance matrix. Each history is presented to each of the two players. Each player tells the referee, for each history, the portfolio which the player believes maximizes EU in (1) for $\lambda = 0.5, 1.0, 2.0$, respectively. The player also provides the referee with the player's own estimate of EU. The referee computes the actual value of EU from the truth, known only to the referee. The referee tabulates the actual value and the players' estimates of this value for the two players. This is repeated for 100 histories per truth and 10 truths for the experiment.

2 The Michaud Player

Michaud proposes the following procedure to handle the fact that observed means, variances, and covariances are not the true parameters but contain noise. In private conversations with the present authors, Michaud points out that more sophisticated procedures could be incorporated into the resampling philosophy. We grant this, but note that it would be difficult to formulate an experiment that encompasses all the possible nuances of both the resampling and Bayesian approaches. The experiment we report here, admittedly, compares "vanilla" resampled frontiers with diffuse Bayes implemented by a particular Monte Carlo analysis.

Following Michaud (1998), the "Michaud player" in our experiment proceeds as follows: given a specific history O ("O" for "Observation") generated by the referee with its means, variances, and covariances, the Michaud player draws 500 new samples of returns on the eight asset classes for 216 months, drawing these from a joint normally distributed i.i.d. random process with the same means, variances, and covariances as O. For each of these 500 samples the Michaud player generates an efficient frontier and then averages these 500 efficient frontiers. Specifically, it notes the first, second, third.... 101st points on the frontier spaced by equal increments of standard deviation. The first point is the one with the highest expected return on the frontier; the 101st point is the one with the lowest standard deviation. The "resampled frontier" has as its first portfolio the average holdings of the first portfolios of the 500 particular frontiers, its second portfolio is the average holdings of 500 second portfolios, etc.

The portfolio mean and variance ascribed to each of the 101 portfolios of the resampled frontier are computed using the original means, variances, and covariances of the observation O. (The present authors thank Richard and Robert Michaud for clarification on this point.) The task that each of the players is assigned is to provide portfolios which maximize the expected value of (1). Therefore, for a given history the Michaud player picks from his resampled frontier the points which maximize the expected value of its estimated EU for $\lambda = 0.5, 1.0$, and 2.0. This process is repeated for each of the 100 randomly drawn histories for each of the 10 truths presented to the player by the referee.

3 The diffuse Bayes player

3.1 Basics

At any moment in time (say $t = 0$) the Bayesian rational decision maker (RDM) acts as if it ascribes a probability distribution $P_0(h)$ to hypotheses h in some space H of possible hypotheses. In the present discussion, a hypothesis is a vector of eight means and 36 distinct variances and covariances:

$$h' = (\mu_1^h, \ldots, \mu_8^h, \sigma_{11}^h, \sigma_{12}^h, \ldots, \sigma_{88}^h) \quad (2)$$

plus the assertion that the variables

$$r' = (r_1, \ldots, r_8) \quad (3)$$

are joint normally distributed with these parameters. The hypothesis space H may be taken as all possible values of h:

$$H = R^{44} \quad (4)$$

It is inconsequential whether we restrict H to the set H^* of 44-tuples that can possibly be parameters of a joint normal distribution, or define it as in (4) and understand that

$$P_0(R^{44} - H^*) = 0 \quad (5)$$

The probability distribution $P_t(H)$ changes over time, as we review below. We assume that, as of any time t, the RDM chooses an action α so as to maximize a single-period utility function

$$EU = \underset{h}{E}[E(U(r;\alpha)|h)] \quad (6)$$

In other words, the action α is chosen so as to maximize EU where U depends on returns r and action α, and the expected return in (6) is computed as if Nature randomly drew a hypothesis h using probability distribution P_t, then drew r given h. In the present experiment the action α is the choice of a portfolio.[2]

To pick a portfolio which maximizes EU in (6) using the utility function in (1), the RDM uses only its estimated portfolio mean (E) and portfolio variance (V) which depend only on its estimated means μ_i of securities and the covariances σ_{ij} (including variances $V_i = \sigma_{ii}$) between pairs of securities. These are given by

$$\begin{aligned}\mu_i &= E(r_i) = \underset{h}{E}(E(r_i)|h) \\ &= \underset{h}{E}\,\mu_i^h \\ &= \text{Avg}\,\mu_i^h \end{aligned} \quad (7)$$

$$\begin{aligned}\sigma_{ij} &= E(r_i - \mu_i)(r_j - \mu_j) \\ &= \text{E}(r_i - \mu_i^h + \mu_i^h - \mu_i) \\ &\quad \times (r_j - \mu_j^h + \mu_j^h - \mu_j) \\ &= \text{E}(\sigma_{ij}^h) - E(\mu_i^h - \mu_i)(\mu_j^h - \mu_j) \\ &= \text{Avg}\,\sigma_{ij}^h - \text{cov}(\mu_i^h, \mu_j^h)\end{aligned} \quad (8)$$

since, e.g.

$$\underset{h}{E}[(r_i - \mu_i^h)(\mu_j^h - \mu_j)|h] = 0$$

In particular, for $i = j$ (Eq. 8) says

$$V_i = \text{Avg}\,V_i^h - \text{Var}\,(\mu_i^h) \quad (9)$$

The last line of (7) and (8) are mnemonics for the immediately preceding lines. These formulas tell us that, for the Bayesian RDM, the expected value of r_i at time t is the average, using $P_t(h)$ over $h \in H$, of μ_i^h; whereas covariance between r_i and r_j is the average σ_{ij}^h plus the covariance between μ_i^h and μ_j^h. In particular, the variance of r_i is the average V_i^h *plus* the variance of μ_i^h.

As evidence accumulates, $P_t(h)$ changes over time, according to the Bayes rule. If $P_t(H)$ has a probability density function $p_t(h)$, and O is an observation taken between t and $t+1$ (e.g. O is the set of monthly returns r_{it} for $i = 1, \ldots, 8, t = 1$ to 216 as described before), with $L(O|h)$ the probability density of O given hypothesis h, then,

$$p_{t+1}(h) = \frac{p_t(h)\, L(O|h)}{\int_H p_t(h)\, L(O|h)\, \mathrm{d}h} \tag{10}$$

The human decision maker (HDM) who wishes to emulate an RDM sometimes avoids the burden of specifying $p_t(h)$ by assuming that

$$p_t(h) = 1/\mathrm{vol}(\Omega^*) \quad \text{for all } h \in \Omega^* \tag{11}$$

where "vol" stands for volume and $\Omega^* \subset H$ is assumed to be sufficiently large that

$$\int_{H-\Omega^*} p_t(h)\, L(O|h)\, \mathrm{d}h \tag{12}$$

is negligible. With (11) assumed, the updated beliefs of (10) become

$$p_{t+1}(h) = L(O|h)/D \tag{13}$$

where

$$D = \int_{\Omega^*} L(O|h)\, \mathrm{d}h \tag{14}$$

and the expected value [with respect to $P_{t+1}(h)$] of any integrable function $v(h)$ is

$$\underset{H}{E}\, v(h) = N/D \tag{15}$$

where

$$N = \int_{\Omega^*} v(h)\, L(O|h)\, \mathrm{d}h$$

In principle, (15) can be used to compute $\underset{H}{E}\, \mu_i^h, \underset{H}{E}\, \sigma_{ij}^h$, and $\underset{H}{E}\, \mu_i^h \mu_j^h$, which are necessary to compute Avg μ_i^h, Avg σ_{ij}^h and cov (μ_i^h, μ_j^h) in (7) and (8). The practical problem is that N and D are integrals over 44-dimensional spaces. As is often done, we will use Monte Carlo analysis to approximate a high-dimensional integral. The specifics of how we do this are described in a following subsection. First, we discuss the fact that a hypothesis space can often be parameterized in different ways, and present a parameterization of the present situation that will be very convenient for the Monte Carlo analysis that follows.

3.2 *Diffuse priors*

Suppose, for the moment, that there was only one unknown parameter, an expected value μ of one random variable r. Then, the standard diffuse prior spreads probability belief concerning μ uniformly over some large interval:

$$p(\mu) = \begin{cases} \dfrac{1}{2\Delta} & \text{for } \mu \in [-\Delta, \Delta] \\ 0 & \text{otherwise} \end{cases} \tag{16}$$

The choice of Δ is not important as long as Δ is sufficiently large, since the contribution to $E(r)$ becomes negligible beyond a sufficiently large Δ. Admittedly, this is often not a very plausible prior. For example, if r is the return on an asset class it is not plausible for the asset class to have a large constant-through-time negative expected return. Such an asset class would disappear. However, the use of (16) is justified as convenient because it saves making a decision as to the exact form to be used for prior beliefs. In effect, it assumes that posterior beliefs are proportional to the likelihood function $L(O|h)$. One justification for assuming posterior beliefs are proportional to $L(O|h)$ is the Edwards *et al.* (1963) principle of stable estimation. "To ignore the departures from uniformity, it suffices that your actual prior density change gently in the region favored by the data and not itself too strongly favor some other region" (p. 202). In particular, it suffices if the likelihood function is much more concentrated than the prior beliefs are, and prior beliefs do not strongly favor any region.

Next, suppose that there are two parameters to be estimated, namely an expected return μ and a standard deviation σ. Now, there are competing choices for a diffuse prior such as

$$p(\mu,\sigma) = \begin{cases} \dfrac{1}{2\Delta_1\Delta_2} & \text{for } \mu \in [-\Delta_1, \Delta_1] \\ & \quad \sigma \in [0, \Delta_2] \\ 0 & \text{otherwise} \end{cases} \tag{17}$$

$$p(\mu, V) = \begin{cases} \dfrac{1}{2\Delta_1\Delta_2} & \text{for } \mu \in [-\Delta_1, \Delta_1] \\ & \quad V \in [0, \Delta_2] \\ 0 & \text{otherwise} \end{cases} \tag{18}$$

$$p(\mu, \log\sigma) = \begin{cases} \dfrac{1}{4\Delta_1\Delta_2} & \text{for } \mu \in [-\Delta_1, \Delta_1] \\ & \log\sigma \in [-\Delta_2, \Delta_2] \\ 0 & \text{otherwise} \end{cases} \tag{19}$$

Since $\log\sigma = \frac{1}{2}\log V$, a similar expression for $p(\mu, \log V)$ would not be a new alternative. Since the use of (16) is justified by convenience and the principle of stable estimation, even when not plausible, one should be permitted the choice between (17), (18), and (19) on the basis of convenience, since the principle of stable estimation would seem to apply about equally to any of them.

With two normally distributed random variables, $r = (r_1, r_2)$, the hypothesis space would most naturally include the choice of

$$h' = (\mu_1, \mu_2, \sigma_1, \sigma_2, \sigma_{12}, \text{or } \rho_{12}).$$

One way of forming diffuse priors for the above is to assume that μ_1, σ_1 and μ_2, σ_2 each have as priors (17), (18), or (19) and that ρ_{12} is independently drawn with a prior density of

$$p(\rho) = \begin{cases} \frac{1}{2} & \text{for } -1 \le \rho \le 1 \\ 0 & \text{otherwise} \end{cases} \tag{20}$$

It might seem that one could repeat the process for $n = 8$ with $\mu_i, \sigma_i, i = 1, \ldots, 8$ having priors (17), (18), or (19) and with each ρ_{ij} independently having (20) as a prior for $i = 1, \ldots, 7$, $j = i+1, \ldots, 8$. One problem with this is that it assigns positive probabilities to correlation matrixes which are logically impossible. For example, it is impossible to have $\rho_{ij} < -\frac{1}{7}$ for every $i \neq j$ for eight returns.

We use a different "diffuse approach" which avoids the above difficulty and is computationally quite convenient for the Monte Carlo analysis described below. This approach uses priors equivalent to nature drawing r_{it} according to

$$p(r_{it}) = \begin{cases} \dfrac{1}{2\Delta} & \text{for } r_{it} \in [-\Delta, \Delta] \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

independently for $i = 1, \ldots, 8$, $t = 1, \ldots, 216$, then computing μ_i, σ_i, and ρ_{ij} as the means, standard deviations, and the correlations of the randomly drawn r_i. The distribution of $(\mu_1, \mu_2, \ldots, \sigma_{88})$ is implicit. In other words, we will find it most convenient to assume that prior probability distribution of $(\mu_1, \ldots, \sigma_{88})$ is the same as that of the sample statistics of random variables $r_1, \ldots, r_8$ drawn uniformly and independently, for sample size $T = 216$. For example, for a very large Δ in (21) the distribution of μ_1 is approximately normally distributed with a large standard deviation.

3.3 Importance sampling

Let

$$K = R^8 \times R^{216} \tag{22}$$

be the space of 8×216 real matrices. Examples of members of K include O, the historical observation handed to each player, and $k_1, \ldots, k_{500}$, the 500 histories which the Michaud player generates. Recall, H is defined in (4) as R^{44}. Members of H include h in (2), the parameters of a joint normal distribution of $(r_1, \ldots, r_8)$.

Let f_{KH} be a function $f_{KH} : K \rightarrow H$ which associates with each point $k \in K$ the $(\mu_1, \ldots, \sigma_{88})$ vector $h \in H$ obtained by computing these parameters from the returns matrix k. For two points k_1, and k_2 in K we define

$$L(k_1|k_2) = L[k_1|f_{KH}(k_2)] = \prod_{t=1}^{216} N[r^t; f_{KH}(k_2)] \qquad (23)$$

where $N(r; h)$ is the normal density of the random vector r given the parameters h. In other words, (23) defines the likelihood of k_1 given k_2 to mean the likelihood of getting the sample k_1 from a normal distribution with parameters $f_{KH}(k_2)$.

Let

$$K^* = \{r \in K| \quad |r_{it}| \leq \Delta \forall i, t\} \qquad (24)$$

for some large Δ. We assume that the prior density is uniformly distributed over this set, K^*. To evaluate an expected value as in (15) by integration would require integration over a large rectangle in an 8×216-dimensional space. This is not feasible. On the other hand, an estimate of $E(v)$ by Monte Carlo, for randomly drawn v, depends on sample size and the moments of v rather than on the dimensionality of K.

Given any function $v(k)$ of the sample point k, in principle, one could estimate the Bayes player's $E(v)$ given O, by sampling k from K^* with probability

$$p(k) = L(O|k)/D \qquad (25a)$$

where

$$D = \int_{K^*} L(O|k)\mathrm{d}k \qquad (25b)$$

Instead, we will have the Bayes player use the same 500 samples from K which the Michaud player uses to compute its resampled frontier. We must keep in mind that these 500 samples were drawn with probability density

$$q(k) = L(k|O) \qquad (26)$$

That is, the 8×216 matrices (r_{it}) in Michaud's samples are drawn joint normally assuming the parameters of the "historical" observation O. Observe that $L(k|O)$ in (26) is not the same as $L(O|k)/D$ in (25a).

There is a standard correction applicable when we wish to estimate an expected value

$$E(v) = \int_{K^*} p(k)\, v(k)\, \mathrm{d}k \qquad (27)$$

and we draw a sample, e.g. $v_1, \ldots, v_{500}$, with probability $q(k)$ rather than $p(k)$. The sample average

$$v^* = \frac{1}{500} \sum_{i=1}^{500} v_i \qquad (28)$$

has expected value

$$E(v^*) = \int_{K^*} q(k)\, v(k)\, \mathrm{d}k \qquad (29)$$

which may differ from $E(v)$ in (27). Instead, we may use a weighted average

$$\bar{v} = \frac{1}{500} \sum \left[p(k)/q(k)\right] v(k_i) \qquad (30)$$

This has expected value

$$E(\bar{v}) = \int q(k) \left[p(k)/q(k)\right] v(k)\, \mathrm{d}k = E(v) \qquad (31)$$

as desired.

The weights in (30) correct for sample probabilities $q(k)$ provided $q(k) > 0$ when $p(k) > 0$. This does not mean that all sampling distributions $q(k)$ are equally good. Since all are adjusted to have the correct $E(v)$, the variance of v depends only on

$$E(v^2) = \int q(k)\left(\frac{p(k)\cdot v(k)}{q(k)}\right)^2 \mathrm{d}k$$
$$= \int \frac{[p(k)v(k)]^2}{q(k)}\mathrm{d}k \tag{32}$$

The minimum of this subject to

$$\int q(k)\mathrm{d}k = 1.0 \tag{33}$$

is

$$q(k) = p(k)|v(k)| \tag{34}$$

Since our sample will serve to estimate the expected value of many different $v(k)$, perhaps all we can conclude from (34) is that it is best to avoid large $q(k)$ where $p(k)$ is relatively small. $q(k) = p(k)$ seems at least good and perhaps ideal.

To compute $p(k)$ we need D from (25b). We now discuss how we approximate this. Let "Vol" be the volume of K^* in (24). Assuming $L(O|k)$may be ignored in $K - K^*$, Eq. (25b) may be written as

$$D = \mathrm{Vol}\int_{K^*}\frac{L(O|k)\mathrm{d}k}{\mathrm{Vol}} \tag{35}$$

This is the volume of K^* times the expected value of $L(O|k)$ in K^* when k is drawn uniformly from it; i.e. with probability density

$$p(k) = 1/\mathrm{Vol} \tag{36}$$

We can approximate this expected value using the sample of 500 k_i drawn with probability density $L(k_i|O)$ by weighing the observed $L(O|k_i)$ by the ratio of desired density (1/Vol) to the density used $L(k_i|O)$. That is, we can approximate the integral (expected value) in (35) by

$$\overline{L} = \frac{1}{500}\sum\frac{L(O|k_i)}{L(k_i|O)\cdot\mathrm{Vol}} \tag{37}$$

But D in (35) is "Vol" times the integral, so the approximation to D is

$$\overline{D} = \frac{1}{500}\sum\frac{L(O|k_i)}{L(k_i|O)} \tag{38}$$

i.e. the observed average ratio of $L(O|k_i)$ to $L(k_i|O)$. It may be objected that if $\overline{D}$ is substituted for D in (25a), $p(k)$ is a ratio of unbiased estimators, which is not necessarily unbiased. On the other hand, with $\overline{D}$ thus used for D in (25a), the weights in (30) sum to one. In this case $\bar{v}$ is a weighted average of $v(k)$, which seems attractive.

Our procedure then is as follows. To approximate the RDMs posterior mean vector μ and covariance matrix C we approximate the expected values of variables v, such as Er_i and Er_ir_j for all i, j, then use the relationships in (7) and (8). To estimate $E(v)$, we evaluate v for each of the Michaud samples from O, namely $k_1, \ldots, k_{500}$, then form the weighted average $\bar{v}$ of the $v(k_i)$ where the weights are shown in (30) with $q(k)$ defined in (26) and $p(k)$ defined in (25a) and (38).

For the one case we checked, most weights $q(k)/p(k)$ are close to unity. Table 2 shows the deciles of the 500 weights computed for Truth 1 History 1. All weights were greater than 0.80 and not greater than 1.025. Ninety percent of the weights were between 0.96 and 1.025. This says that the Michaud sample is a good one for the present

Table 2 Distribution of weights $p(k)/q(k)$ for Truth 1 History 1.

Deciles	From	To
1st	0.809	0.961
2nd	0.961	0.984
3rd	0.984	0.999
4th	0.999	1.005
5th	1.005	1.012
6th	1.012	1.015
7th	1.015	1.019
8th	1.019	1.021
9th	1.021	1.023
10th	1.023	1.025

purpose, according to (34) and the discussion that follows it.

4 Results

The results of the experiment are presented in Tables 3 and 4. The first panel of Table 3 shows averages of estimated and actual expected utility achieved by the two players. Specifically, for $\lambda = 0.5, 1.0$, and 2.0, as indicated by the row labeled "Lambda", and for each player, as indicated by the row labeled "Player", the table presents two columns of information. The first column is the average (over the 100 histories generated for a truth) of the players' estimate of expected utility. The second column is the average of actual utility as evaluated by the referee. For example, on the line labeled Truth 1 we see that, on average over the 100 histories generated for Truth 1, the Bayesian player believed it had achieved an expected utility of 0.01181 whereas the average of its actual *EU* was 0.00712. The comparable numbers for the Michaud player are 0.01032 and 0.00753. Thus, both players overestimated how well they did, but the Michaud player overestimated less and achieved more. On the next nine lines similar numbers are reported for Truth 2 through Truth 10. The final three lines of the panel summarize results for the 10 truths. In particular, the average over the 10 truths of the Bayesian player's estimate was 0.01383 but it actually achieved 0.00861. In the average over all 10 truths, again, the Michaud player overestimated less and achieved more. In fact, comparing the average *EU* each player achieved in each of the 10 truths, the average over the 100 histories was greater for the Michaud player than the Bayes player in the case of each of the 10 truths, as noted in the last line of Panel A.

A similar story holds for $\lambda = 1.0$ and 2.0. Looking at the last row of Panel A for the actual *EU* achieved by the two players for these cases we see that the Michaud player achieved a higher average (over the 100 histories for a given Truth) in 10 out of 10 truths for $\lambda = 1.0$ and 2.0.

For some individual histories of the 100 histories of a given truth, the Bayes player had a higher *EU* than the Michaud player. In fact, in Panel B of Table 1 the entry for $\lambda = 0.5$, Bayes player, Truth 1 reports that the Bayes player achieved a higher *EU* than the Michaud player in 54 out of the 100 histories, despite having a lower average over the 100. Sticking with Truth 1, the Bayes player also "won" 54 out of 100 times for $\lambda = 1.0$, and 50 out of 100 for $\lambda = 2.0$. The Bayes player's "win count" was even more favorable in the case of Truth 6. In this case, the Bayes player "beat" the Michaud player 62 times out of 100 for $\lambda = 0.5$, 60 for $\lambda = 1.0$ and 66 for $\lambda = 2.0$. Nevertheless, the average *EU* achieved, averaged over the 100 histories, was higher for the Michaud player in each of these Truths.

Panel C of Table 1 shows that, for a given Truth, the standard deviation of the achieved *EU* was higher for the Bayesian than the Michaud player. For example, for Truth 1, $\lambda = 0.5$, the standard deviation of *EU* for the Bayesian player was 0.00210 as compared to 0.00132 for the Michaud player. In fact, for all three values of λ and all 10 truths, the variance of the actual *EU* was lower for the Michaud player than the Bayes player.

Most significant for our purpose is the fact that the Michaud strategy delivered higher average *EU* in 10 out of 10 truths for three out of three values of λ. Thus, the Michaud player did a better job of achieving the objective, namely high *EU*.

Table 4 displays the results of a slightly different game. In this second game, for each history and each truth each player computes an efficient frontier as in the first game. But instead of picking a point from the frontier for each λ, the player passes its entire frontier to the referee. For each λ the referee picks the point on the player's frontier that has the

Table 3 Player's choice of portfolio.

λ:	0.5	0.5	0.5	0.5	1.0	1.0	1.0	1.0	2.0	2.0	2.0	2.0
Player:	Bayes	Bayes	Michaud	Michaud	Bayes	Bayes	Michaud	Michaud	Bayes	Bayes	Michaud	Michaud
Eval. by:	Player	Referee	Player	Referee	Player	Referee	Player	Referee	Player	Referee	Player	Referee
Panel A: EU *averaged over 100 histories, for each of 10 truths*												
Truth 1	0.01181	0.00712	0.01032	0.00753	0.01004	0.00564	0.00886	0.00594	0.00754	0.00394	0.00678	0.00426
Truth 2	0.01528	0.00885	0.01194	0.00901	0.01389	0.00783	0.01085	0.00801	0.01160	0.00616	0.00902	0.00664
Truth 3	0.01011	0.00614	0.01009	0.00737	0.00887	0.00534	0.00904	0.00636	0.00692	0.00410	0.00721	0.00481
Truth 4	0.01457	0.00850	0.01147	0.00862	0.01324	0.00746	0.01041	0.00763	0.01094	0.00573	0.00849	0.00600
Truth 5	0.01170	0.00641	0.00984	0.00694	0.00988	0.00480	0.00846	0.00549	0.00706	0.00282	0.00612	0.00322
Truth 6	0.01646	0.01056	0.01304	0.01078	0.01462	0.00890	0.01149	0.00914	0.01173	0.00670	0.00911	0.00700
Truth 7	0.01590	0.01147	0.01408	0.01152	0.01412	0.00989	0.01271	0.01015	0.01124	0.00758	0.01036	0.00793
Truth 8	0.01502	0.00956	0.01261	0.01005	0.01329	0.00811	0.01119	0.00861	0.01053	0.00578	0.00866	0.00610
Truth 9	0.01402	0.00906	0.01241	0.00961	0.01204	0.00719	0.01087	0.00798	0.00892	0.00462	0.00812	0.00521
Truth 10	0.01343	0.00846	0.01130	0.00900	0.01176	0.00676	0.00975	0.00735	0.00909	0.00402	0.00712	0.00453
Grand mean	0.01383	0.00861	0.01171	0.00904	0.01217	0.00719	0.01036	0.00767	0.00956	0.00514	0.00810	0.00557
Std Dev	0.00205	0.00171	0.00138	0.00150	0.00200	0.00161	0.00133	0.00145	0.00190	0.00148	0.00129	0.00142
No. times better		0		10		0		10		0		10
Panel B: Number of "wins" out of 100 histories, for each of 10 truths												
Truth 1		54		46		54		46		50		50
Truth 2		52		48		54		46		65		35
Truth 3		46		54		43		57		41		59
Truth 4		57		43		61		39		64		36
Truth 5		43		57		27		73		30		70
Truth 6		62		38		60		40		66		34
Truth 7		57		43		53		47		42		58
Truth 8		54		46		48		52		41		59
Truth 9		32		68		28		72		27		73
Truth 10		61		39		49		51		52		48
Avg No. wins		51.80		48.20		47.70		52.30		47.80		52.20
No. times better		7		3		5		5		5		5

Panel C: Standard deviation of EU *over 100 histories, for each of 10 truths*

Truth 1	0.00516	0.00210	0.00401	0.00132	0.00470	0.00160	0.00359	0.00102	0.00356	0.00116	0.00265	0.00077
Truth 2	0.00445	0.00149	0.00395	0.00084	0.00416	0.00185	0.00372	0.00113	0.00364	0.00258	0.00331	0.00121
Truth 3	0.00354	0.00244	0.00339	0.00109	0.00331	0.00231	0.00321	0.00092	0.00286	0.00178	0.00284	0.00080
Truth 4	0.00413	0.00203	0.00369	0.00101	0.00398	0.00221	0.00356	0.00118	0.00377	0.00235	0.00331	0.00117
Truth 5	0.00514	0.00161	0.00347	0.00077	0.00475	0.00126	0.00329	0.00064	0.00396	0.00077	0.00275	0.00058
Truth 6	0.00469	0.00223	0.00476	0.00114	0.00427	0.00227	0.00438	0.00101	0.00379	0.00242	0.00373	0.00117
Truth 7	0.00544	0.00178	0.00347	0.00088	0.00515	0.00155	0.00331	0.00086	0.00447	0.00111	0.00296	0.00073
Truth 8	0.00399	0.00217	0.00413	0.00119	0.00391	0.00169	0.00398	0.00112	0.00376	0.00117	0.00360	0.00073
Truth 9	0.00584	0.00144	0.00371	0.00056	0.00551	0.00130	0.00359	0.00070	0.00467	0.00089	0.00319	0.00065
Truth 10	0.00399	0.00283	0.00445	0.00155	0.00391	0.00212	0.00422	0.00130	0.00360	0.00166	0.00353	0.00077
Avg Std Dev		0.00201		0.00104		0.00182		0.00099		0.00159		0.00086
No. times better		0		10		0		10		0		10

Table 4 Referee's choice of portfolio.

λ:	0.5	0.5	1	1	2	2
Player:	Bayes	Michaud	Bayes	Michaud	Bayes	Michaud
Eval. by:	Referee	Referee	Referee	Referee	Referee	Referee
Panel A: EU *averaged over 100 histories, for 10 truths*						
Truth 1	0.007811	0.007709	0.006303	0.006253	0.004852	0.004899
Truth 2	0.009625	0.009407	0.008641	0.008594	0.006967	0.007104
Truth 3	0.007111	0.007552	0.006198	0.006647	0.004741	0.005139
Truth 4	0.009157	0.008915	0.008109	0.008049	0.006220	0.006395
Truth 5	0.006721	0.007008	0.005200	0.005662	0.003504	0.003661
Truth 6	0.011378	0.011183	0.009781	0.009608	0.007486	0.007481
Truth 7	0.011935	0.011571	0.010425	0.010303	0.008178	0.008260
Truth 8	0.009674	0.010071	0.008225	0.008714	0.005799	0.006309
Truth 9	0.009423	0.009641	0.007712	0.008169	0.005112	0.005576
Truth 10	0.008193	0.008854	0.006665	0.007339	0.004151	0.004718
Grand mean	0.009103	0.009191	0.007726	0.007934	0.005701	0.005954
Std Dev	0.001701	0.001509	0.001654	0.001473	0.001506	0.001414
No. times better	5	5	5	5	1	9
Panel B: Number of wins out of 100 histories, for 10 truths						
Truth 1	65		60		53	
Truth 2	66		64		58	
Truth 3	56		52		47	
Truth 4	69		67		59	
Truth 5	52		32		40	
Truth 6	67		60		60	
Truth 7	74		64		54	
Truth 8	44		45		37	
Truth 9	55		43		27	
Truth 10	58		58		41	
Avg wins	60.6		54.5		47.6	
No. times greater	9		7		5	
Panel C: Std Dev of EU *over 100 histories, for 10 truths*						
Truth 1	0.00169	0.00112	0.00113	0.00082	0.00063	0.00043
Truth 2	0.00110	0.00058	0.00128	0.00060	0.00127	0.00061
Truth 3	0.00186	0.00083	0.00163	0.00071	0.00117	0.00060
Truth 4	0.00143	0.00062	0.00153	0.00060	0.00139	0.00056
Truth 5	0.00145	0.00053	0.00106	0.00042	0.00054	0.00027
Truth 6	0.00106	0.00086	0.00074	0.00058	0.00099	0.00059
Truth 7	0.00131	0.00084	0.00102	0.00070	0.00090	0.00054

Table 4 (*continued*)

λ:	0.5	0.5	1	1	2	2
Player:	Bayes	Michaud	Bayes	Michaud	Bayes	Michaud
Eval. by:	Referee	Referee	Referee	Referee	Referee	Referee
Truth 8	0.00137	0.00089	0.00121	0.00076	0.00093	0.00051
Truth 9	0.00137	0.00053	0.00125	0.00049	0.00082	0.00039
Truth 10	0.00274	0.00145	0.00223	0.00114	0.00120	0.00067
Avg Std Dev	0.00154	0.00082	0.00131	0.00068	0.00098	0.00052
No. times lower	0	10	0	10	0	10

highest true *EU*. Game 2 thus addresses the question of whether the superiority of the Michaud player over the diffuse Bayesian player in the first game is due to a better frontier or to a better pick from an equally good frontier.

The Bayes player does much better in Game 2 than it did in Game 1. In particular, for $\lambda = 0.5$ and 1.0 Panel A of Table 4 shows that with five out of 10 truths the Bayesian player achieves higher average *EU* than the Michaud player as compared to 0 out of 10 in Game 1. Also, Panel B shows that for $\lambda = 0.5$ and 1.0 the Bayesian player has a higher *EU* in many more histories for a given Truth than the Michaud player. On the other hand, the Michaud player comes out ahead overall. In particular, for every λ the "Grand Mean" of achieved *EU* averaged over all truths is greater for the Michaud player than the Bayesian player. However, the outperformance of the Michaud player over the Bayes player is smaller in the second game than in the first. In particular, for $\lambda = 0.5$ the difference in performance between the two players is only about 20% as great in the second game as it is in the first ($0.000088 = 0.009191 - 0.009103$ versus $0.00043 = 0.00904 - 0.00861$), about 44% as great when $\lambda = 1.0$ and 59% as great when $\lambda = 2.0$.

As explained in the next section, for $\lambda = 0.5$, *EU* in (1) is approximately[3] $E(\ln(1+r))$. This, in turn, is $\ln(1+g)$ where g is the geometric mean or growth rate. We can, therefore, give the results in Tables 3 and 4 a more concrete interpretation for the case of $\lambda = \frac{1}{2}$. Annualizing, the Bayes player believes it can achieve an "average"[4] annual growth rate of 18.05% ($0.180548 = \exp(12 \cdot (0.01383)) - 1$), whereas the portfolios it chose had an average actual growth rate of 10.89% and the best from its frontier averaged a growth rate of 11.54%. The Michaud player thought it could achieve an average annual growth of 15.09%; the portfolios it chose had an average growth rate of 11.46%; the actual average highest growth portfolio on its frontier was 11.66%. Thus, in game 1, the Michaud methodology adds 0.57 to the average growth rate. In game 2 it adds 0.12.

The relatively better performance of the Bayesian player in Game 2 (as compared to its performance in Game 1) suggests that the Game 1 superiority of the Michaud player is more due to a wise pick from its frontier than due to a superior frontier, though the latter reason is also applicable.

5 Questions

The preceding results raise questions for portfolio theory and practice. In particular, the results represent something of a crisis for the theoretical foundations of portfolio theory as presented in Part IV of Markowitz (1959), Chapters 10–13. Chapters 10 through 12 present introductory accounts of utility analysis as justified by Von Neumann and Morgenstern (1944), personal probability as justified by Savage (1954), and dynamic programming as presented by Bellman (1957). Chapter 13 applies these principles to the problem of selecting a portfolio. Specifically, mean–variance analysis is justified as an approximation to the single-period "derived" utility function always associated with many-period utility maximization. It is argued that the mean–variance approximation should be good as long as the probability distribution of return is not spread out too much. Calculations—by Markowitz (1959), Young and Trent (1969), Levy and Markowitz (1979), Dexter *et al.* (1980), Pulley (1981, 1983), Kroll *et al.* (1984), Simaan (1987) and Hlawitschka (1994)—show that, for most utility functions proposed for practice, the mean–variance approximation to expected utility is quite robust. As Levy and Markowitz conclude

> If Mr. X can carefully pick the E,V efficient portfolio which is best for him then Mr. X, who still does not know his current utility function, has nevertheless selected a portfolio with maximum or almost maximum expected utility.

In addition, Markowitz and van Dijk (2003) illustrate the ability of a suitably constructed "single-period" mean–variance analysis to give near-optimum results in the case of transaction costs and changing probability distributions. One caveat however: as Grauer (1986) illustrates, the return distributions from highly levered portfolios are too spread out for mean–variance approximations to do well. However, for unlevered return distributions as considered in the present paper, computations have generally shown mean–variance to be quite good.

Thus, until now, calculations seem to support the theoretical foundations for mean–variance analysis presented in Part IV of Markowitz (1959). An integral part of these foundations is that a RDM will use probability beliefs where objective probabilities are not known, and will update these beliefs according to the Bayes rule as evidence accumulates. Usually, when Bayesian inference is tried in practice it is assumed that, prior to the sample in hand, beliefs are "diffuse"—i.e. "neutral" in some sense with respect to which hypothesis is true—as recommended by Jefferies (1948) or Edwards *et al.* (1963).

Given this background, the results presented in this paper are badly in need of an explanation. Such explanation could be in terms of why Bayesian updating did not do better, or why the Michaud estimation did so well.

Concerning why Bayesian updating did not do better: it may have to do with the difference between the computation which we performed and which a RDM would perform. The latter is an integration over a high-dimensional space, well beyond foreseeable human computational abilities. We approximated this integral by Monte Carlo sampling. (Note the distinction between the sample which the referee handed both players, and the sample we used to approximately compute the integral which the RDM computes exactly.) If this—exact versus approximate calculation of updated beliefs—is the source of difficulty with the Bayesian approach taken here, then, maybe the conclusion will be that Bayesian inference is ideal for the RDM but not for the human, at least at the level of computational effort spent by the Bayesian and Michaud players in the reported experiment.

Alternatively, perhaps the problem with the approach taken here is the priors used. Perhaps "diffuse prior" should be defined differently. Or, perhaps, an informed prior should be used like those of Black and Litterman (1990)—but updating the priors using history rather than user estimates as in Black and Litterman.[5]

Expected utility and Bayesian inference were originally proposed, by Daniel Bernoulli and Thomas Bayes in the Eighteenth Century, as plausible rules for action when the future is unknown (see Bernoulli, 1954; Bayes, 1958). Von Neumann and Morgenstern (1944) and Savage (1954) derive these rules from more basic principles of rational behavior. The resampled frontier as presented by Michaud (1998) is a plausible procedure which, we find, works quite well. But how does it relate to the theory of rational behavior? Does it contradict one or more of Savage's axioms? If so, is this a black mark against the method or against the axioms? Or does Michaud's procedure somehow satisfy the Savage axioms? We would very much like to know the answers to some or all of these questions.

Practical questions, raised by the success of the Michaud method in the experiments reported here, include those of costs and benefits. In particular, how much expected return do these procedures add for a given level of risk—in practice. This may involve transaction costs, changing probability distributions, non-normal distributions—all assumed away in the current experiments. Historical backtests might shed some light on these matters.

Concerning costs, computation costs may or may not be a problem. It does not take long or cost much these days to generate a set of 500 frontiers and average these. But it might still be computationally burdensome to compute many such resampled frontiers in a backtest with many monthly re-optimizations, with the backtest frequently repeated to see the effects of alternate parameter settings. However, a Bayesian update of beliefs would also be computationally burdensome in such a case.

Finally, the cost of using a resampled efficient frontier depends on what the patent holder charges for the use of this patented procedure (see note 1).

6 Conclusions

This paper reports the results of an experiment comparing two procedures for dealing with sampling error in the inputs to a mean–variance analysis. One procedure is the Bayesian updating of diffuse priors. The other is Michaud's resampled efficient frontier. In the experiment a referee generates 10 "truths" at random from a "seed" distribution. From each "truth" the referee randomly generates 100 histories. Each history is presented to a Bayesian player and a Michaud player. Each player follows its prescribed procedure to determine which portfolio would provide highest $E - \lambda V$ for $\lambda = 0.5, 1.0$, and 2.0. Sometimes one player, sometimes the other picks a portfolio with higher $E - \lambda V$. But in the case of each truth and each value of λ, the average of the 100 values of $E - \lambda V$ is higher for the Michaud player than the Bayes player. However, the Bayes player does almost as well as the Michaud player when each player presents its entire efficient frontier to the referee, and the referee picks the player's best portfolio from the frontier. This suggests that the chief problem with the Bayesian player's choice of portfolio is that the latter is more over-optimistic than is the Michaud player in estimating achievable portfolio mean and variance.

This result has practical implications for the estimation of inputs to a mean–variance analysis, even for methods other than the two considered explicitly here. For example, in practice, mean–variance analysis is often performed at an asset class level with estimates of means based partly on judgment, but using historical variances and covariances. The results of this paper imply that these variance estimates are too low. First, if you accept the theory of rational behavior under uncertainty developed by

Savage (1954), as explained by Markowitz (1959) Chapter 12, then you should not use historical variance, nor even an average variance—averaged over possible explanations of history. Rather, you should use the latter *plus* a term reflecting your uncertainty in your estimate of the mean. Furthermore, the results of the present paper imply that, for reasons unknown to us, when this theoretical correction is made, the investor is still too optimistic for his or her own best interest.

Acknowledgments

The authors thank Anthony Tessitore for extremely valuable suggestions.

Notes

1 Resampled efficiency, as described in Michaud (1998, Chapters 6 and 7), was co-invented by Richard Michaud and Robert Michaud and is a US patented procedure, #6,003,018, December 1999, patent pending worldwide. New Frontier Advisors, LLC, has exclusive licensing rights worldwide.

2 The assumption of a single-period utility function is not less general than the assumption of many-period or continuous-time utility maximization, since many-period or continuous-time utility maximization may be reduced to a series of one-period or instantaneous utility maximizations using a "derived" utility function, as described by Bellman (1957). In general, the time-varying derived utility function U_t may be a complicated function that includes state variables as well as returns, and depends on what has gone before. Our specific assumption, that U_t is given by (1), is a vast simplification which we justify on the grounds that our objective is not to solve the dynamic programming problem for some many-period or continuous-time investment model, but to take a reading on the ability of two alternate methods to handle uncertainty.

3 For other values of λ, EU in (1) approximates the expected value of other utility functions. The choices made by a Bernoulli/Von Neumann and Morgenstern utility function are not affected by adding a constant or multiplying by a positive constant. That is, the same decisions maximize $E[a + bU(r)]$, $b > 0$, as those that maximize $EU(r)$. It is, therefore, essential to the validity of the comparisons made in the text—e.g. that the difference in performance is only 20% as great in game 2 as game 1 when $\lambda = 0.5$—that this comparison is in fact unaffected by the arbitrary choice of a and $b > 0$.

4 The "average" referred to here is the antilog of an average logarithm, therefore, a geometric mean.

5 Harvey *et al.* (2003) reports the results of an experiment in which Bayes outperforms Michaud when conjugate priors are used.

References

Bayes, T. (1958). "Essay Toward Solving a Problem in the Doctrine of Chances: with a biographical note by G. A. Barnard." *Biometrika* **45**, 293–315. (Also published separately by the Biometrika Office, University College, London.) *Philosophical Transactions of the Royal Society* 370–418, 1763.

Bellman, R. E. (1957). *Dynamic Programming*. Princeton, New Jersey: Princeton University Press.

Bernoulli, D. (1954). "Specimen theoriae novae de mensura sortis. Exposition of a New Theory on the Measurement of Risk" (English translation by Louise Sommer). *Econometrica* **22**, 23–26. (Originally published in 1738. *Comm. Acad. Sci. Imp. Petropolitanae* **5**, 175–192.)

Black, F. and Litterman, R. (1990). "Asset Allocation: Combining Investor Views with Market Equilibrium." *Journal of Fixed Income* Goldman Sachs, September.

Dexter, A. S., Yu, J. N. W., and Ziemba, W. T. (1980). "Portfolio Selection in a Lognormal Market When the Investor Has a Power Utility Function: Computational Results." In: M. A. H. Dempster (ed.), *Stochastic Programming*. New York: Academic Press, 507–523.

Edwards, W., Lindman, H., and Savage, L. J. (1963). "Bayesian Statistical Inference for Psychological Research." *Psychological Review* **70**(3), 193–242.

Grauer, R. R. (1986). "Normality, Solvency, and Portfolio Choice." *Journal of Financial and Quantitative Analysis* **21**(3), 265–278.

Harvey, C. R., Liechty, J. C., Leichty, M. W. and Muller, P. (2003). Portfolio Selection with Higher Moments, working paper, Duke University, Durham, NC.

Hlawitschka, W. (1994). "The Empirical Nature of Taylor-Series Approximations to Expected Utility." *The American Economic Review* **84**(3), 713–719.

Jeffreys, H. (1948). *Theory of Probability*. Oxford: Clarendon Press.

Kroll, Y., Levy, H., and Markowitz, H. M. (1984). "Mean Variance Versus Direct Utility Maximization." *Journal of Finance* **39**(1) 47–61.

Levy, H. and Markowitz, H. M. (1979). "Approximating Expected Utility by a Function of Mean and Variance." *American Economic Review* **69**(3), 308–317.

Markowitz, H. M. (1959). *Portfolio Selection: Efficient Diversification of Investments*. John Wiley & Sons New York: 2nd edn., Basil Blackwell, Cambridge, MA.

Markowitz, H. M. and Usmen, N. (1996a). "The Likelihood of Various Stock Market Return Distributions, Part 1: Principles of Inference." *Journal of Risk and Uncertainty* **13**, 207–219.

Markowitz, H. M. and Usmen, N. (1996b). "The Likelihood of Various Stock Market Return Distributions, Part 2: Empirical Results." *Journal of Risk and Uncertainty* **13**, 221–247.

Markowitz, H. M. and van Dijk, E. L. (2003). "Single-Period Mean–Variance Analysis in a Changing World." *Financial Analysts Journal* **59**(2), 30–44.

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

Pulley, L. M. (1981). "A General Mean–variance Approximation to Expected Utility for Short Holding Periods." *Journal of Financial and Quantitative Analysis* **16**, 361–373.

Pulley, L. M. (1983). "Mean-Variance Approximations to Expected Logarithmic Utility." *Operations Research* **31**(4), 685–696.

Savage, L. J. (1954). *The Foundations of Statistic*, 2nd edn. Dover Publications, Inc., New York: John Wiley & Sons.

Simaan, Y. (1987). "Portfolio Selection and Capital Asset Pricing for a Class of Non-Spherical Distributions of Assets Returns." PhD Thesis, Baruch College, The City University of New York.

Von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior, Princeton University Press*, 3rd ed., Wiley, 1967.

Young, W. E. and Trent, R. H. (1969). "Geometric Mean Approximation of Individual Security and Portfolio Performance." *Journal of Financial and Quantitative Analysis* **4**, 179–199.

Keywords: Resampled frontier; Bayesian analysis; diffuse Bayes; mean–variance analysis; sampling errors; Michaud.