

May 30, 2008

H. Gifford Fong, Editor
Journal of Investment Management
3638 Mt. Diablo Blvd.
Suite 200
Lafayette, CA 94549

RE: Comments on Bayes vs. Resampling: A Rematch

Dear JOIM Editor:

We appreciate the opportunity to respond to comments by Richard and Robert Michaud, on our recent article “Bayes vs. Resampling: A Rematch”. While we agree that there are some overlapping elements to the Bayesian approach that we put forward and the Resampling approach promoted by the Michauds, there are some substantial differences that merit closer attention. The main difference is that the Resampling approach breaks from the traditional optimization framework of maximizing an expected utility and instead takes the expectation of weights that maximum a utility; stated simply the Resampling approach maximizes and then averages instead of maximizing an average (or an expected return). This represents a fundamental departure from the seminal framework proposed originally by Markowitz (1959).

Although we are criticized for not following the Michauds’ recommendations regarding implementation of the Resampling methodology, something that we address in a post script at the end of this letter, we wish to focus our discussion on the differences in optimization approaches and within that framework discuss the role of various investment scenarios (referees) that could be used to assess the performance of an asset allocation strategy. Due to the limited nature of this format we will of course limit our discussion to simple examples.

There are three components to both of the approaches being considered: 1) generation of random parameters, 2) the optimization framework used to determine an optimal set of investment weights and 3) the investment scenario used to determine how well the resulting weights perform. The first issue is important, but not of real interest as the Monte Carlo approach used in the Resampling methodology can be viewed as an approximation to the MCMC sampler used to generate posterior draws of the mean and covariance matrix.

The Resampled Optimization Approach

The second point is of substantial interest as this point represents a major break from the traditional (or more accurately characterized, the dominate) approach to optimal decision making. Part of the challenge with past discussions of the Resampled optimization approach, is that these discussions have been restricted to finding weights for long only portfolios (*i.e.* the weights are constrained to be positive). If we lift this restriction, then

we can derive analytic results which help clarify the differences between the Resampled optimization approach and the Traditional optimization approach. To illustrate, in a way that is directly comparable with the Resampling approach, we will show differences using the assumption that investor's utilities are a function of parameter values.

Traditionally, an investor will choose weights, w , that maximize their expected utility $u(w, \mu, \Sigma)$, or assuming quadratic utility they would solve the following problem:

$$w_T = \arg \max E[u(w, \mu, \Sigma) | H_o] = \arg \max (w^T E[\mu | H_o] - \lambda w^T E[\Sigma | H_o] w)$$

where the subscript 'T' denotes the traditional approach and 'R' denotes the Resampled approach, H_o represents the observed history, $E[\mu | H_o]$ are the expected returns, $E[\Sigma | H_o]$ is the variance-covariance matrix, and λ is the risk aversion of the investor; using simple calculus we can obtain the standard set of optimal weights

$$w_T = \frac{E[\Sigma | H_o]^{-1} E[\mu | H_o]}{2\lambda}. \quad (1)$$

An investor, who is following the Resampling approach, inverts the traditional order and they use the following weights

$$w_R = \frac{1}{N} \sum_{i=1}^N w_i^*(\mu_i, \Sigma_i)$$

where

$$w_i^* = \arg \max (w^T \mu_i - \lambda w^T \Sigma_i w) = \frac{\Sigma_i^{-1} \mu_i}{2\lambda}, \quad \mu_i, \Sigma_i \sim f(\mu, \Sigma | H_o),$$

where $f(\mu, \Sigma | H_o)$ represents the parameter uncertainty given the observed history H_o . A direct comparison of the weights that result from the two approaches is instructive:

$$w_T = \frac{1}{2\lambda} \left(\frac{1}{N} \sum_{i=1}^N \Sigma_i \right)^{-1} \left(\frac{1}{N} \sum_{i=1}^N \mu_i \right) = \frac{1}{2\lambda} \overline{\Sigma}^{-1} \overline{\mu} \quad \text{and} \quad w_R = \frac{1}{N} \sum_{i=1}^N \frac{\Sigma_i^{-1} \mu_i}{2\lambda} = \frac{1}{2\lambda} \overline{\Sigma^{-1} \mu},$$

where $\overline{\Sigma}^{-1} \overline{\mu}$ is product of the averages, and $\overline{\Sigma^{-1} \mu}$ is the average of the products.

To help understand the differences consider the single asset case where an investor can either invest in a single risky asset or in the risk free asset. In this case, the Resampled investor's weight will always be larger, in absolute value, than the traditional investor, or $\Pr(|w_R| > |w_T|) = 1$; this is a simple consequence of Jensen's Inequality. The intuition for this can be seen from the fact that μ and Σ are essentially uncorrelated and that small values of Σ , when inverted, will have an increasingly larger impact than large values of Σ . This can be seen by recalling that $(1/\Sigma)$ is a hyperbola and as a

result $(1/\Sigma) \xrightarrow{\Sigma \rightarrow 0} \infty$. Hence, averaging over $(1/\Sigma)$ will result in a larger value than taking one over the average of Σ .¹

Selecting the Referee

Since we have derived the weights, we can explicitly calculate the expected utility for both investors and compare their performance, if we can determine an appropriate investment scenario or stated differently if we can agree on an acceptable referee. The key to understanding the referee's perspective is to recall that the investor creates their weights, w_T and w_R , from moments based on the observed history H_O (e.g.

$\bar{\mu} = E[\mu | H_O]$) and the referee can use a different history, H_{REF} , to evaluate the performance of the weights. For example if we continue with our single asset example, the expected utility becomes

$$EU_T = E[u(w_T, \mu, \Sigma) | H_{REF}] = \frac{1}{2\lambda} \left(\bar{\Sigma}^{-1} \bar{\mu} E[\mu | H_{REF}] - \frac{(\bar{\Sigma}^{-1} \bar{\mu})^2 E[\Sigma | H_{REF}]}{2} \right)$$

and

$$EU_R = E[u(w_R, \mu, \Sigma) | H_{REF}] = \frac{1}{2\lambda} \left(\overline{\Sigma^{-1} \mu} E[\mu | H_{REF}] - \frac{(\overline{\Sigma^{-1} \mu})^2 E[\Sigma | H_{REF}]}{2} \right).$$

Using simple algebra, we can explicitly determine when the traditional approach will have a higher expected utility; $EU_T > EU_R$ when

$$E[\Sigma | H_{REF}]^{-1} E[\mu | H_{REF}] < \frac{1}{2} (\bar{\Sigma}^{-1} \bar{\mu} + \overline{\Sigma^{-1} \mu}) = \bar{\Sigma}^{-1} \bar{\mu} + \Delta, \quad (2)$$

where $\Delta = \frac{1}{2} (\overline{\Sigma^{-1} \mu} - \bar{\Sigma}^{-1} \bar{\mu})$, $\bar{\mu} = E[\mu | H_O]$, $\bar{\Sigma} = E[\Sigma | H_O]$, and as noted above,

$\Pr(\Delta > 0) = 1$.

There are three referees that we would like to consider the Predictive (or One Step Ahead) Referee, the Truth Referee, and the Random Referee. The Predictive Referee uses the observed history, $H_{REF} = H_O$, or the predictive distribution based on the observed history (which is the only history available to the investor), hence

$$E[\mu | H_{REF}] = E[\mu | H_O] = \bar{\mu} \quad \text{and} \quad E[\Sigma | H_{REF}] = E[\Sigma | H_O] = \bar{\Sigma}.$$

Based on (2), traditional approach will always win, with probability 1. We are willing to concede that the Predictive Referee is backwards looking, in that he/she uses just the

¹ Jensen's Inequality simply says that for a convex function, the function of the average will be less than the average of the function.

observed history to assess performance, but we are not willing to concede that this is a circular argument. It simply points out that the traditional approach always beats the Resampled approach, when the objective function being maximized is used to assess performance.

The Truth Referee uses the parameters, $\mu_{True}, \Sigma_{True}$, that were used to generate the history, which is equivalent to using an infinite history, or

$$H_{REF} = H_{\infty} = \{r_{\tau}, \tau = 1, \dots, \infty : r_{\tau} \sim f(r | \mu_{True}, \Sigma_{True})\}.$$

Hence, the resulting moments are the true parameters, or $E[\mu | H_{\infty}] = \mu_{True}$ and $E[\Sigma | H_{\infty}] = \Sigma_{True}$. While the Truth Referee is often assumed to be the best referee, he/she has the fatal flaw that the investor must hold the portfolio forever; only at that point will there be no variability in the parameter estimates and only then will the investor's utility agree with the utility used by the Truth Referee. We do not find this infinite time horizon scenario to be a creditable scenario, even if other feels this is viable.

The Random Referee acknowledges the shortcoming of the Predictive and Truth Referee and assumes that the investor will hold the portfolio for a finite amount of time, e.g. the amount of time equal to the original history. Over this future time a new history based on the true parameters $\mu_{True}, \Sigma_{True}$, will be generated, or

$$H_{REF} = H_{Rand} = \{r_{\tau}, \tau = 1, \dots, \tau_{Rand} : r_{\tau} \sim f(r | \mu_{True}, \Sigma_{True})\} \text{ and } H_{Rand} \neq H_O.$$

For the Random Referee, the resulting moments, $\mu_{Rand} = E[\mu | H_{Rand}] \cong \bar{\mu}_{Rand}$ and $\Sigma_{Rand} = E[\Sigma | H_{Rand}] \cong \bar{\Sigma}_{Rand}$, will be used to assess the weights. It is worth noting that in the limit, as the size of the new history goes to infinity $\tau_{Rand} \rightarrow \infty$, the Random Referee becomes the Truth Referee.

The resulting empirical moments, $\bar{\mu}_{Rand}, \bar{\Sigma}_{Rand}$, can be viewed as random variables drawn from the same distribution as the original empirical moments, $\bar{\mu}, \bar{\Sigma}$, or

$$(\bar{\mu}_{Rand}, \bar{\Sigma}_{Rand}), (\bar{\mu}, \bar{\Sigma}) \sim f(\bar{\mu}, \bar{\Sigma} | \mu_T, \Sigma_T).$$

This means that the Traditional approach will do better than the Resampled approach more than 50% of the time. To see this recall that $EU_T > EU_R$ when

$$E[\Sigma | H_{REF}]^{-1} E[\mu | H_{REF}] = \bar{\Sigma}_{Rand}^{-1} \bar{\mu}_{Rand} < \bar{\Sigma}^{-1} \bar{\mu} + \Delta$$

and realize that and both $\bar{\Sigma}_{Rand}^{-1} \bar{\mu}_{Rand}$ and $\bar{\Sigma}^{-1} \bar{\mu}$ have the same distribution. A simple symmetry argument requires that if $\Delta > 0$, then

$$\Pr(\bar{\Sigma}_{Rand}^{-1} \bar{\mu}_{Rand} - \bar{\Sigma}^{-1} \bar{\mu} < \Delta) > 0.5 \quad \text{or} \quad \Pr(EU_T > EU_R) > 0.5.$$

Finally, because $\Pr(\Delta > 0) = 1$, due to Jensen's Inequality, the Random Referee will more times than not, declare the Traditional optimization approach to be better than the Resampling optimization approach; which is not surprising given the overwhelming acceptance of the traditional decision science definition that an optimal decision is one that maximizes expected utility.

Sincerely,

Campbell Harvey, John Liechty and Merrill Liechty

P.S.

With regards to specific criticism relating to how we implemented the Resampled algorithm in our recent paper "Bayes vs. Resampling: A Rematch," Harvey, Liechty and Liechty (HLL). We readily acknowledge that we used a variation of the Resampled approach used by in Markowitz and Usmen (MU) in their 2003 paper (also published in this journal), but in doing so we followed recommendations explicitly given by the proponents of the Resampled approach. For example, we were criticized for using the λ -associated Resampling algorithm instead of the rank-ordered algorithm. When the λ -associated method is put forth in Michaud's (1998) book they say, "As a practical matter, the choice between the two approaches may simply be a matter of convenience." (See appendix of chapter six, from Michaud 1998, page 67).

We also agree that the issue of sample size is important. HLL and MU follow the recommendations given by Michaud (1998) for the Resampled method which says that 500 Monte Carlo samples should be used to find the Resampled Efficiency portfolio. We agree that the reassessment of this guideline given by Michaud in their open letter, *i.e.* that 500 samples is probably too small, is probably correct, but the reader should recall that this is the particular recommendation that was given previously by Michaud for this specific set of data. For a proper Bayesian analysis, 500 samples are too restrictive. Therefore HLL uses the recommended number of samples from each discipline, 500 for the Resampled method and (something much bigger) 25,000 for the Bayesian method. We would like to note that it is not our recommendation to integrate over 44 dimensions with a sample size of 500. Finally, we would like to remind the reader that our primary purpose in HLL was not to determine whether the guidelines put forward for the Resampled approach were optimal; we find it ironic that both of the main criticisms of the HLL comparison of the Resampled approach are that HLL followed Michaud's own recommendations for implementing the Resampled approach.