NEW FRONTIER

A Note on Statistics vs. Finance: Is Return Easier to Estimate than Risk?

By

David N. Esch, Ph.D.

New Frontier Advisors, LLC
Boston, MA 02110

July 2013

Investment practitioners often claim that very stable estimates are available for standard deviations and correlations of assets. In contrast, they maintain that the expected returns of those same assets have relatively unreliable estimates available. Indeed, many software and investment firms have built large practices selling risk models. While far from above critiques, such estimates are generally respected in the community and widely used. However, the purported ease of estimating risk compared to expected return is contrary to provable facts from statistics which dictate that for any dataset and analysis, population moments are estimated with expanding error distributions as the order of the moment increases. This implies that variance estimates have a wider range of error than mean estimates for the same dataset. This simple mathematical fact remains widely ignored by the investment community. How is it possible that this mathematical fact is somehow false in asset management?

It is important to understand the different perspectives of a professional statistician from an investment practitioner. When practitioners talk about estimating the mean return of securities, they are typically talking about something other than the statistician estimating a global static mean from sample dataset. Markets are constantly churning, in an equilibration process that is responding to various stresses and shocks. So the true means, variances, and return distributions evolve moment to moment. However, every instant in time corresponds to only one market price per security, so it is impossible to precisely estimate these non-static means and variances, separating "signal" from "noise." In order to produce useful estimates, practitioners estimate something simpler than dynamically changing parameters, such as a static mean and variance over an interval of time. Reducing the number of estimated quantities enables more precise estimation of a mean vector and variance matrix via various estimation techniques, the simplest being the sample mean and variance of the data stream. This mismatch of the simplified estimation model with a known complexity of markets is the key to explaining the reversal in available precision of the first two central moments: mean and variance.

The sample variance is generally not a bad mismatch to the shifting variance about the shifting mean that we really want to estimate. In theory, the total variance could be decomposed into the sum of the true variance about the shifting mean, and another variance, the variance of the shifting mean about its theoretical long-term center, the ergodic[1] mean. Since the variance we're really interested in, the former, is likely to be the dominant component of the total variance, it is reasonable to substitute the static variance estimate for the one we want at the

[1] The ergodic mean is a theoretical center of the probability distribution of a quantity over all time. Since return distributions shift over time, expected returns for investable assets at any specific time are likely to differ from the corresponding ergodic means.

terminal end of the estimation window. This estimate probably captures correlations among estimates rather well. The shifting variances of the market may wobble around these values, but this wobbling is rather inconsequential in any type of portfolio allocation analysis compared to the uncertainty in the means. Since variances and correlations are only really meaningful across multiple observations anyway, the fluctuations of the dynamic variance are second-order considerations and getting a precise estimate for a longer time period is desirable. For larger analyses risk factors have been well researched and do an excellent job of capturing the individual and joint forces affecting asset price movement.

On the other hand, there is a major mismatch between the shifting mean we're interested in (the expectation at the final instant of the estimation window) and the static mean of the entire time period. The latter, while estimable, is actually a reasonable estimate for the ergodic[1] mean of the asset returns, which has little investment value. When managers think about the mean input to an optimizer, they are thinking of alphas, i. e. expectations calculated from external information and likely relevant to the current time window only. These types of mean inputs are far more useful in creating a wealth-generating portfolio, for returns fluctuate, often rather vigorously, around their long-term means but can exhibit short-term tendencies or biases over business cycles that, if captured in the alphas, can lead to enhanced portfolio performance. When price movements settle down in low volatility periods, there is no guarantee that they land in any particular place. The tendency for mean reversion in expected returns for most assets is weak at best and many return series can be well characterized by a random walk model confined to some reasonable global limitations on values of returns.

The above reasons form a good justification for using methods such as James-Stein or hierarchical linear models (HLM), for instance, that reduce noise and shrink historical data when estimating historical means, and adding as much external information as possible through methods like subjective prior distributions created from investor views in a Bayesian analysis. These external views can come from general theories of market behavior such as risk premia, from specific current news trends, or from security-specific information. Aggregating securities into group estimates, as the Bayesian analyses do, creates more stable estimates by using more observations within the pooled groups to add stability missing from the one security at a time analysis.

Most human analysts lack the necessary intuition to construct a reliable covariance matrix without computer assistance. Furthermore, moving correlation coefficients are difficult to imagine and return series behavior is well explained by the static historical model. For these reasons, the NFA Asset Allocation System by default uses a historical covariance matrix, even when views are present which greatly affect the mean estimate.

The term “mean” return has multiple interpretations in finance, while it is unambiguous within a particular statistical model. The finance practitioner's most specific notion of "mean" is likely to correspond to a statistical model that is not estimable from historical data. There are numerous examples in finance of statistical or mathematical terminology having multiple interpretations, especially as the statistical models increase in complexity and require more precise and/or arcane terminology. It is helpful to understand these distinctions when designing estimation models for optimizers, and to recognize that in order for any model to be useful it must simplify some of the complexity of the system since we don’t have the luxury of replicating the market to get more data. We must make do with the information we have. This mandates adding simplifying structure to our estimation models to add strength of multiple observations to the estimates, and to take advantage of as many information sources as possible when designing portfolios to maximize expected wealth and reduce risk. Exactly how to make these compromises is where intuition meets rigorous analysis in quantitative finance, and indeed in statistics as a whole, and an important part of what separates a skilled quantitative analyst from a less skilled one.