Greek Alphabet Soup and Risk-Adjusted Performance

by Arun S. Muralidhar

In the last year, every journal or magazine seemed to carry yet another piece on "alpha" versus "beta" or the next new performance measure termed "kappa" or "omega". Some measures of performance go by the names of the inventors (unfortunately, not Greek) like the Sharpe ratio, Sortino ratio or M-square (Modigliani and Modigliani), Morningstar ratio or those of friends/family of the authors like the SHARAD (my measure which stands for Skill, History and Risk-Adjusted Performance – however, this is also my brother's name).

In this note, we aim to help clients understand some of these measures and add a few Greek alphabets that seem to be missing in the debate, but which are important. Moreover, we will attempt to lay down some basic principles that a risk-adjusted performance measure should meet (based on theoretical and practical considerations) before the author can claim that this is the ultimate measure, replacing all others.

Arun S. Muralidhar is currently Chairman of Mcube Investment Technologies, LLC and Managing Director of FX Concepts, Inc. Prior to this he was Managing Director and Head of Currency Research, JP Morgan Investment Management. Prior to JP Morgan, he was a member of the World Bank's Investment Management Committee which oversees more than $40 billion in assets, including assets of the staff pension fund

Alpha and Beta, but where art Rho? Is sigma significant?

The new paradigm espoused by all is "separate alpha from beta" and "do not pay alpha fees for beta," but what does all this mean? Let us start from the perspective of a pension fund (the same would apply to an endowment, or insurance company or even hedge fund FoF), understand how they function and then put these terms into context by relating back to finance theory from where they came.

Typically, a Trustee Committee sets objectives for the fund and determines a long-term strategic asset allocation (SAA) for the fund. This SAA is typically of the form: 60% in Stocks (measured by the S&P500 index), and 4030% in Bonds (measured by the Lehman Aggregate Bond Index) and 10% in Alternatives (with a lot of murkiness as to whether the index is LIBOR + 5%, S&P 500 + 3% or even some HFRI-type index). I simplify greatly as there are typically many more asset classes and more stratification, but this will suffice for our purposes.

This benchmark is given to the investment team at a pension fund (which unfortunately many trustees do not trust entirely and hence consultants are hired) who are expected to implement this portfolio and ideally achieve the basic benchmark return and then some. However, trustees often do not want the staff to take too much risk relative to the SAA.

Staff have two potential activities to beat the benchmark – make all investment decisions internally (very rare) or structure the portfolio optimally (e.g., either overweight/underweight stocks, keep some assets in cash or assign a manager the Russell 3000 stock index as a benchmark and manage the risk between the S&P500 and the Russell 3000) and hire external managers to implement these mandates.

We will focus on the latter, but again the problem is that the staff do not entirely trust that the managers are skillful and this is where the Greeks appear. The same would apply to pension funds that hire Hedge Fund FoFs, who then hire hedge funds managers. In this case, there is a potential doubling of delegation and more scope for "mistrust".

The concept of beta has its origins in the Capital Asset Pricing Model, where the beta is a measure of covariance of an asset with that of the market portfolio. In investing parlance, very often when clients and asset managers refer to beta, they typically mean either the SAA or the market index to which they are measured (many of which would fail the required assumptions of the market portfolio of the CAPM). The two concepts of market portfolio and beta are completely different, but one would not know that if they read all the marketing pitches today about beta. Moreover, the true beta that a pension fund requires is relative to its liabilities – something the CAPM is completely silent about. Hence a lot of the discussion on "beta" is wrong for the average pension fund.

Alpha is also very loosely defined in investing parlance – many people use alpha to mean returns in excess of the benchmark index returns with no information at all about how correlated (or "rho") these returns are to the benchmark (either SAA or market index). The correlation coefficient is very important because, ideally, to lower the risk of achieving a target rate of return the investor should find as many uncorrelated return (or ideally negatively correlated positive return) streams as possible. So the more appropriate definition of alpha, and largely an outgrowth of the hedge fund industry, is to consider the return streams that are uncorrelated with the SAA or market index. A good example of such a return stream in a classic pension portfolio is the pure excess return/cash generated from active currency strategy or from a "market neutral" hedge fund.

However, in the traditional mandates given out by pension funds, the manager was asked to beat the benchmark index and because the clients did not know exactly what the manager was doing back in their offices, they tried to ring fence their activities relative to the benchmark by monitoring their tracking error (or the volatility of excess returns). In the FoF case, this is compounded, though most investors would claim they care about absolute risk and not relative risk. There are three elements that go into a tracking error calculation – the volatility of the market index (or sigma of the index), the volatility of the external manager's portfolio (or sigma of the manager) and the correlation of the manager's performance with the benchmark index (or another rho).

For many years, investors and rating agencies looked solely at the excess return relative to a market index and divided them by some measure of risk to establish whether managers were good or bad. In some cases, the measure of risk was the tracking error (creating the information ratio or modified-Sharpe ratio), or a measure of volatility of only underperformance (creating the Sortino ratio), or alternative measures creating the Morningstar ratio (Gambera 2004). It took a seminal article, which it appears has not been widely read by researchers on risk-adjusted performance, by the late Franco Modigliani and Leah Modigliani (Modigliani and Modigliani 1997) to start to bring the focus back on sigma and also making measures of risk-adjusted performance become more user friendly. Their measure has many interesting implications for the Hedge Fund industry, with the focus on absolute risk.

Towards a higher standard for risk-adjusted performance measures

They pointed out that when offered a choice of investment alternatives relative to a benchmark (and the benchmark could even be LIBOR for hedge funds), investors should ideally look for two things: (a) a risk-adjusted performance measure that is expressed in percentage terms (i.e., return terms as that is what benefits can be paid from); and (b) help the investor construct better portfolios without getting hung up on the jargon. I use their starting point and my experience on different sides of the investment table to lay out a series of criteria that researchers should try to achieve in developing new risk-adjusted performance measures. These criteria are meant to help bridge the gap between theoretical rigor and practical value.

i) Expressed in the metric relevant to the investor

The late Franco Modigliani used to argue that getting a risk-adjusted performance measure as a ratio did nothing for the client and could often lead them to incorrect decisions. If the client could create measures of risk-adjusted performance that were

expressed in return terms, then they could think about whether this return would be helpful to pay pension benefits or tuitions for endowments. As Modigliani-Modigliani in the *M-square* measure demonstrate, an information ratio of –0.5 (manager A) may be better than an information ratio of 0.3 (manager B). It all boils down to a simple condition – if the volatility (sigma) of the manager A's performance is very low and that of manager B very high, comparing returns is an apples-to-oranges comparison.

The returns relative to the benchmark cannot be compared until one normalizes for the differences in the sigmas. When one implements their M-square transformation, they show very credibly that a negative information ratio manager could be preferred in risk-adjusted return terms to the positive information ratio manager. We show why in the next section. The Omega ratio gives clients a measure of the performance profile relative to some performance threshold, but the M-square is cleaner as it is in percent terms.

ii) Help with portfolio construction

One of the other problems with ratios is what to do with them? For example, what does one do with a Sharpe ratio of 0.5 or an information ratio of 0.7 or an Omega of 1.2? Modigliani and Modigliani (1997) take a neat extension of the Sharpe measure to create risk-adjusted portfolios out of an investment possibility. In short, they used the risk-free asset (cash, assumed to have zero volatility) to lever or delever the active manager's portfolio.

For example, if manager A (with a negative information ratio) took too little sigma risk, the investor can borrow cash to lever the manager allocation so that the sigma of the risk-adjusted portfolio (or as they called it, the RAP) is equal to that of the benchmark. Similarly, manager B may need to be delevered. Therefore, the additional value of their measure was it told the investor how much to allocate to different active strategies and how to balance that with allocations to cash, giving useful advice on whether leverage can help meet a performance threshold.

However, if the investor's measure of risk is the tracking error as opposed to sigma (i.e., worried about relative risk rather than absolute volatility), then these two risk-adjusted portfolios may have the same sigma, but very different correlations to the benchmark. In other words, the two tracking errors could be different and hence the M-cube measure of performance was created to normalize for both the differences in volatilities and correlations (to achieve a target tracking error). The M-cube is an extension of the M-square and now gives the investor allocation information as to how to allocate to cash (risk-free asset); the benchmark (market portfolio or the asset with zero tracking error or "beta") and the potential active strategy. Notice that none of the newer measures of performance for hedge fund strategies like the Kappa or the Omega (Kazemi at al 2004, Kaplan and Knowles 2004) or other ratios can offer this helpful advice to investors.

The value of this approach is that clients can hire a manager and without altering their strategy create appropriate exposures to "beta" and "cash" and thereby achieve the optimal portfolio rather than looking only for "alpha", which appears to be a rare commodity – all the marketing pitches notwithstanding.

iii) Incorporate the length of history or skill of the manager

One of the shortcomings of all measures is that there is no time element to them and this is particularly vexing in the hedge fund world. Is an Omega over 2 years of data better than a lower Omega over 5 years of data? The problem investors have is (a) they would like to have some concept of the validity of the statistic (which reduces as the number of data points falls); (b) ideally compare managers with different lengths of data histories without throwing out any data; and (c) say something about the skill of the manager.

One of the issues with the old Sortino ratio was that picking up only observations when the return was less than the benchmark greatly reduced the sample set on which the sigma was calculated. This has been adjusted for in using bootstrap techniques to generate this ratio. It can be shown that including time could make the M-square and M-cube measures incorrect and hence a measure was proposed in Journal of Performance Measurement, called the SHARAD (Skill, History and Risk-Adjusted performance), where time was used to

establish the confidence the investor could have that the manager was skillful and this was then multipled by the risk-adjusted performance measure. This is a good fix but will suffer from the problem of requiring a unique objective function to make such a measure appropriate. We discuss this below in section v.

iv) Consistency

One interesting fact in currency management or CTAs is that the industry wide excess returns of managers and high relative to their tracking error. However, the returns are "lumpy" – i.e., performance is often achieved in a few months of the year. No institutional investor or FoF likes a positive "alpha" strategy with lumpiness as it requires explaining the same or defending it vis-à-vis the Trustee Board or a client.

Two statistics that can help capture the lumpiness of returns are the "hit-rate" (or number of positive periods divided by total number of periods), where anything above 50% is better than a coin toss and drawdown statistics, or how much and how long can you lose relative to a high water mark before you get fired. I like to call the drawdown statistics 'Yield to Fire!" Most traditional risk adjusted performance measures do not tackle consistency and it typically relates to the emotional aspects of investing and delegating duties.

v) Satisfy complex objectives/satisfy higher moments

One of the knocks against the Sharpe ratio was that it assumes that investors have mean-variance preferences and hence did not capture higher moments such as skew and kurtosis. This could lead to potentially bad decisions to invest in strategies that had high blow up risk. The classic example here is a strategy to sell naked options, which gives many small positive returns, but could have one catastrophic negative return that bankrupts the investor, could have a good Sharpe ratio. This is where measures such as the Omega could dominate in capturing higher moments.

However, such a strategy is not a bad one if coupled with another that has the exact opposite profile like a classic trend strategy – many small losses with a few large gains. The combination of the two if properly designed (i.e., where the correlation between them is negative) can give a much better overall portfolio profile for the investor. Therefore, the Omega could penalize a strategy in isolation without giving an investor insight into what other strategy to combine the same with – which the M-square approach articulated very neatly.

vi) Ability to rank individual managers and combinations of managers

One aspect sorely lacking in the discussion is how various measures of manager risk-adjusted performance are helpful when clients hire not one, but multiple managers. Rarely, do FoFs, retail or institutional investors hire a single manager or mutual fund and hence the risk-adjusted performance measures or ranking schemes that cannot incorporate a portfolio approach can be misleading. For example, the Morningstar rating could lead to a retail investor buying a number of individually (highly) rated funds, but it has been shown that this can be a poor decision in the context of overall portfolio construction as the Morningstar technique does not lend itself easily to multiple manager portfolios.

Summary

This note set out to clarify some of the discussion on alpha, beta, rho, sigma, omega, kappa etc. and help highlight criteria by which future research on risk-adjusted performance measures can be directed. The goal of this list, by no means exhaustive, is to find good theoretical underpinnings to these measures yet find consistency with the practicalities of managing portfolios. The practical issues become much more relevant in managing institutional portfolios as there are multiple layers of decision making and with less than complete trust at all layers.

Moreover, investors rarely hire a single manager and this adds an additional layer of complexity. The goal of the investor is to maximize risk-adjusted return for whatever definition of risk, and simple ratios cannot help in this process. What this piece shows is that in the end investors may require more than one measure to satisfy complex objectives, but before the next researcher makes "gyro"ic (pronounced "heroic") claims about the next new greek alphabet performance measure, that they check that they meet these basic criteria. ■

The author thanks the late Franco Modigliani and Sanjay Muralidhar for assistance and encouragement in developing these ideas. These are the personal views of the author and are theoretical in nature. They do not represent the views of Mcube Investment Technologies, FX Concepts, Inc. or any of their affiliates.