

Alpha Estimation and the Definition of Asset Specific Risk

Dan diBartolomeo

Northfield Webinar

January 2015

Alpha Estimation and Specific Risk

- Under the definition of alpha used by many quant managers, the opportunity set available to the investor to obtain extraordinary risk-adjusted returns (i.e. alpha) by security selection is presumed to be a linear function of the security level specific risk magnitudes
 - Grinold (Journal Portfolio Management, 1994)
- The impact of specific risk estimation on alpha estimation and marginal risk contribution jointly impact the optimality of portfolio position sizes.

Alpha is Beauty in the Eye of the Beholder

- Alpha is return in excess of the expectation of the CAPM
 - Jensen (1968)
- Alpha is all return in excess of a specified benchmark
 - Industry definition: could arise from factor tilts, factor timing, security selection
- Alpha is excess return not arising from fixed factor tilts
 - Common definition in academic research (Fama, French, Carhart)
- Alpha is excess return arising from security selection, orthogonal to known risk factor effects
 - Grinold (1994), $\text{Alpha} = \text{IC} * \text{Volatility} * \text{Score}$

The Practical Inconsistency

- Almost all “quant” strategies are based on associating differences in expected returns with some set of numerical attributes which are common across securities
 - If we assume returns are linearly related to factor exposures (i.e. the APT holds) and your risk model has no omitted variables every “alpha factor” has to be a linear combination of risk factors
 - The opportunity set for active management must therefore relate to the manager’s ability to predict factor returns *which is specifically excluded from the Grinold definition*, yet many managers use the Grinold formula for scaling alpha magnitudes
 - diBartolomeo(1998), <http://www.northinfo.com/Documents/56.pdf>
 - Bulsing, Sefton and Scowcroft (2004)

Taking Grinold at Face Value

- If we agree that the Grinold definition of alpha is the right one, estimation of security specific risk is a key element
 - It has been the custom of the asset management industry to use “asset specific risk”, “idiosyncratic risk” and “residual risk” all interchangeably.
 - When we model risk at the security level, we allocate observed variations in security returns to linear relationships with “common factors” that are believed to be relevant in differing degrees to all securities. The factors are thereby responsible for explaining the covariance of securities, including a large portion of the variance of a given security.
 - Such assumptions of linear return/risk relationships are the underpinning of theories of capital market equilibrium such as the CAPM and APT.

Model Estimation

- To estimate the factor model we typically use standard techniques such as regression analysis.
 - In a typical “OLS” regression we fit the relationship between the dependent variable (security returns) and the independent variables (factor exposures or factor returns depending on the type of model) so as to minimize the “residuals”, the portion of the security returns that are not explained by the factors.
 - The standard deviation of the residuals is closely related to the standard error of the regression.
 - To the extent we cannot explain the variation in security returns as coming from a common factor, it is customary and *seemingly logical* to assume that the residual returns arise from causes which are idiosyncratic in nature and specific to the particular asset in question.

The Negative Definition Problem

- The negative specification is conceptually unsatisfying
 - *We are defining asset specific risk based on our assessment of what we don't know, rather than any kind of affirmative definition in which we would try to quantify and explain asset specific risk based on what we do know.*
- Having an affirmative of specific risk may be very useful
 - Better assessment of volatility at the individual security level which is needed for good alpha scaling
 - Better assessment of the marginal contribution of individual security positions to portfolio risk
 - How the expectation of risk and volatility may be different in long term as compared to a short horizon

The Reconciliation Problem

- Another drawback of the negative definition is the inability to easily reconcile risk estimates across different traded securities within the capital structure of a company.
 - The estimate of firm specific risk we might get from residual return volatility in stock returns might be very different from what we would get from observing residual return volatility from the corporate bonds of the same firm.
 - While theoretical constructs exist to link across the capital structure such as Merton (Journal of Finance, 1974), the heterogeneous nature of investors and liquidity levels across the different types of financial instruments makes the analysis complex.

The Time Horizon of Alpha and Specific Risk

- A more affirmative definition of asset specific risk could also help us draw the distinction between “stock specific risk” and “firm specific risk”.
- It seems intuitive to suggest that stock specific risk might include short-term variations in stock returns caused by time-varying investor behavior, which would not necessarily have any impact on the underlying business risks at the firm level.
 - In the long run, transient investor-driven effect should wash out leaving long term estimates of stock specific risk to converge to the firm-specific risk level.

Model Dependent Ambiguity of Magnitude

- The final unpleasant aspect of the negative definition of asset specific risk is *ambiguity of magnitude*.
 - The residual risk of a stock's return is the portion of the observed return variation over some sample period that is not explained by the chosen set of factors
 - The estimated value of asset specific risk will be impacted by the nature and estimation process of the risk model in use.
 - **We will see what we have chosen to see by virtue of many choices:** the sample period over which returns are observed, the frequency of return observations (daily returns have very different properties than monthly returns), any weighting scheme for the observations, the number of factors chosen, the particular nature of each factor chosen to participate in the risk model as an independent variable

Estimation Dependence

- Magnitudes of residual risks will also depend on the particular statistical procedures used to estimate the various components of the model.
 - Some factors such as balance sheet leverage relate to bankruptcy risk and therefore to highly non-linear return events, which may lead to dependence between residual risk magnitudes and factor exposures
 - Another issue is whether the model is based purely on the past, or like Northfield models is conditioned on contemporaneous information in order to differentiate between short-horizon and long-horizon forecasts.
 - For a useful discussion of these issues see diBartolomeo (Chapter 14, [Oxford Handbook of Quantitative Asset Management](#), 2012,).

-

Assumptions of Geometric Brownian Motion

- Portfolio theory assumes that security returns are normally distributed with consistent volatility and no serial correlation.
 - If observed return distributions that do not have these IID properties, we make adjustments in the estimated total volatility of a stock, which impacts the magnitude of residual variance
 - The adjustment is based on a method from Parkinson (Journal of Business, 1980) in which security volatility is mathematically inferred from the “high/low” price range observed for the security during each observation during the sample period.
 - Empirical research papers have shown that “rescaled range” measures of volatility are more efficient than the traditional time series calculation because we are not ignoring what happens between the start and end points in time for each observation .

Is Volatility Decoupled from Fundamentals?

- Govindaraj, Livnat, Savor and Zhao (2013)
 - Presented this conference last year in Montreal
<http://www.northinfo.com/documents/566.pdf>.
 - All occurrences of a one day stock return with absolute value greater than five percent were recorded over a large universe of stocks from 1982 to 2011. During the period from 1982 to 2011, the study then observes changes in analyst's earnings estimates for the week *subsequent* to the price change. From 1999 to 2011, they also observed changes in the analyst's target prices for the stock over the same subsequent week.
 - The premise is that if the large price changes ($\text{abs} > 5\%$) were investor's response to actual changes in the business conditions of the firm, either the earnings estimates or the share price targets would change as well.

Much Ado About Nothing

- If the large price changes of a specific security were the result of “uninformed” investor activity such as a shareholder making a wholesale change in asset allocation from stocks to bonds, it would be unrelated to the firm’s business and therefore we would expect no change in the earnings estimates or share price targets.
 - It was found that no changes in earnings estimates or share price targets occurred in approximately 80% of the events.
 - This lends substantial credence to the idea that the volatility of stock’s returns is greater than the operating risks of the underlying firms. It also suggests that longer term estimates of asset specific risk should be lower, as the random shocks associated with transient investor behavior average out.

A Capital Structure Approach

- Another approach to making an affirmative definition of asset specific risk would be through analysis of the capital structure of a firm.
 - If a firm has debt, the possibility of defaulting on the debt will be approximately coincident with the value of the equity of the firm going to zero, as described by the “contingent claims analysis” method of Merton (1974).
 - If a firm has traded debt with observable market yields, we can subtract out the yield on a comparable duration riskless debt (i.e. Treasuries) to obtain a credit related yield spread.
 - To the extent that corporate bonds are often callable before maturity, the credit spread is usually adjusted into a measure called an “option adjusted spread” (OAS). This is the extra yield that investors demand for the credit risk of that firm.

Further Down the Merton Path

- In diBartolomeo (Journal of Investing, 2010) a Merton style model is presented that links stock volatility to the volatility of the firm's underlying assets, and to the resultant probability of default (PD) on debt when the value of the firm's assets goes below the face value of the firm's debt.
- This study shows that given stock prices and observed stock volatilities for US stocks (including ADRs) from 1992 to 2009, the implied "half-life" of firms was about 20 years, comparable to a default probability of 3.4% per annum.
- To the extent that bond defaults are actually less frequent than 3.4%, we might conclude that stock volatility is upward biased relative to the actual risks of the firm, consistent with our inferences drawn from the results of Govindaraj, et. al.

Specific Risk at the Firm Level

- We can also exploit this process to *estimate specific risk at the individual firm or stock level*.
 - As described in diBartolomeo (2012) we can also use the Merton process to estimate the “loss given default” (LGD) on a corporate bond.
 - The product of PD and LGD is the “expected loss” (EL) due to credit risk. OAS is the lower bound of the market’s estimate of expected loss.
 - If we can observe the OAS on the bonds of a particular firm, we can work backward through the model to solve for the corresponding level of stock volatility (and variance). We can then subtract out the stock return variance explained by our factors and arrive at a bound on the implied estimate of specific risk

Accounting for Non-Normality

- One subtlety associated with obtaining an OAS implied value for specific risk is that a non-zero probability of default will mean that the distribution of equity returns will be non-normal with left skew and kurtosis.
- To incorporate these higher moments into an equivalent volatility measure, we suggest following approximation, derived from Wilcox (Journal of Portfolio Management, 2003)

$$S^* \sim (S^2 - (200/(6*S)) [(2/3)* MS^3 + (1/2)* KS^4])^{.5}$$

S^* = adjusted volatility

S = traditional volatility

M = skew

K = kurtosis

Investors Do Respond to the Real World

- We can also use the fact that investors are responding to information about real world events to help us form an affirmative definition of specific risk.
 - Kyle, Obizhaeva, Sinha and Tuzun (2012) shows that a theoretically predicted relationship between the frequency of news articles on companies, and the volatility of their stocks was fit almost perfectly by the empirical data over hundreds of companies and many years.
 - Another paper is diBartolomeo, Mitra and Mitra (Quantitative Finance, 2009) which showed that the volatility of major company shares in the US and Europe was more effectively predicted by mathematical summarization of news articles than by the standard method of implying volatility of the market prices of traded stock options.

Impact of News

- Since the 2009 publication Northfield has continued research on the relationship between news and volatility, including sponsoring projects for two teams of MIT graduate students.
 - Stock volatility of some firms was consistently related to both the number and “sentiment” (good/bad) of news stories, while no relationship was detected with other firms.
 - The impact of particular events on daily volatility decayed quickly to insignificant levels in a month or less.
 - Daily variation in intra-day volatility is very closely related to news flows (T statistic > 9) on a sample of 1.7 million data points).
- At least for short term forecasts, it is apparent that incorporating the impact of news on asset specific risk estimates is another essential element of an affirmative estimation process.

A Quick Empirical Test of Ambiguity

- It should be obvious that models of risk that are based on different security universes, with different observation frequencies, or estimated over different historic sample periods could produce different estimated values for stock specific risk.
- What is less obvious is the extent to which the specific risk estimates could vary by changing the nature of the risk model while holding all other aspects of the process (universe, sample period, observation frequency) constant.
- To test this proposition we compared the vector of estimated stock specific risks for the Northfield US Fundamental Model and the Northfield US “Single Market” Model as of August 31, 2014.

A Tale of Two Models

- As the name Fundamental model implies, in this model the factors are defined by endogenous characteristics of stocks such as industry group, book/price ratio, market cap and momentum. The statistical process needed to estimate the model parameters is a cross-sectional regression analysis across the universe of stocks done separately for each observation month.
- To the contrary, the US Single Market model uses exogenous information like interest rates, oil prices and the return on the market as factors. The statistical estimation process is a time series process done separately for each stock (like Fama-French).
- Both models use a 60 month sample period, a comparable degree of exponential weight of observations and the same Parkinson adjustments to the total volatility estimate.

Measuring Ambiguity in 2014

- Our results showed that the cross-sectional correlation between two vectors of stock specific risk values was approximately 87% across a universe of about 5000 stocks.
- The distribution of the differences at the individual stock level had an extremely high degree of kurtosis of more than 14, implying that the two models generally agreed but there was large disparity in a small number of cases.
 - In such cases, the magnitude of alpha forecasts would be very sensitive to the choice of risk model
- Comparable tests on the same data for August 2013, and August 2003 produced nearly identical results with cross-sectional correlations even higher at over 90%, and a kurtosis level over 12.

A Two Way Estimation Compromise

- It may also be possible to generate convergence between the estimation methods.
 - We can certainly take the endogenous model, estimate factor returns for each time period and then use the time series of factor returns in a time series regression. This would produce new estimates of previously observed factor exposures.
 - We would then repeat the cross-sectional estimation to get new factor returns and cycle through the process again. Eventually, the process would converge to stable values.
 - Similarly, we can take the time series factor exposures as given, and use cross-sectional estimation for factor returns. The new factor return times series would replace the original values in the times series regressions. Again, iterating through the process would eventually reach convergence

Conclusions

- There is an inherent conflict between the widely used Grinold alpha scaling formula and the strategies to which it is applied.
 - If we accept the Grinold formula, then accurate estimation of asset specific risks is vital.
- We are not doomed to having a “negative” definition of asset specific risk.
 - There are clear means by which we can distinguish between stock specific risk and firm specific risk.
 - The variation in specific risk over time seems well explained by quantification of news flows.