June 2014

Northfield News

A Newsletter for the Friends and Clients of Northfield

Definition and Decomposition of Risk of Investment Portfolios

By Dan diBartolomeo

A vast amount of time and effort is expended in trying to understand the risks of investing in financial markets. Despite this plethora of effort, we assert that a large fraction of investment practitioners have at best a poor understanding of risk and the economic appropriateness of particular metrics of risk for their sort of organization. This article will first provide an overview of the conceptual basis of risk, and then discuss numerous, often ambiguous aspects of current practice.

Risk Definition

In terms of the matter at issue, we start from three basic principles. The first principle is that the purpose of investment in risky assets is to generate a return above the risk free rate rather than as an exercise to minimize asset risk. If the overriding purpose was to minimize risk, investors could simply hold cash or a zero coupon bond that provided the appropriate dollar pay off for a given future consumption expenditure. The second principle is that investors do not assess risk as an abstract exercise but rather in the context of a decision making process which explicitly or implicitly represents their utility function. This second requirement suggests that the choice of risk measure will arise out of the nature of that utility function. The final principle is that the current discussion will address *risk* only in terms of the potential of an undesirable event arising from a known distribution of possible outcomes. To the extent that in the real world, the distribution of possible future outcomes is only estimated with some degree of error, we will reserve *parameter uncertainty* for a separate discussion.

Most financial institutions operate along the lines of two broad utility functions. The first goes back to the original work of Nicholas Bernoulli (1713) and Daniel Bernoulli (1738) on stochastic dominance. Put simply, investors prefer more return to less, they don't like risk (but have a hard time defining what risk is) and there is a decreasing marginal utility to wealth (if I have enough money to buy 20 houses, I really don't care very much about getting more wealth to buy house 21). The simplest function that fits these ideas is to assume that investors are seeking to maximize the expectation of the log of their wealth. This has been the most widely used assumption for economic analysis of investment for many decades. To the extent that you have a log-wealth utility function, the use of variance as the risk measure arises naturally as analytically derived in Markowitz and Levy (1979).

The Markowitz-Levy paradigm has been extended in a couple pieces of key literature. The first to which I would draw your attention is Wilcox (2000) that incorporates both higher moments and time-varying risk aversion based on the balance sheet capacity of the investor to withstand loss. The Wilcox process is a more precise framework for dynamic risk mitigation strategies such as Constant Proportion Portfolio Insurance (CPPI). It should also be noted that recent research such as Steve Ross's paper on his "Recovery Theorem" (2013) provides evidence of investor behavior consistent with Wil-

(Definition, Continued on page 4)

Special Points of Interest:

- Main Article: Definition and Decomposition of Risk of Investment Portfolios
- Annual Conference Announcement: Stowe Vermont
- Tech Tip: Analytical Considerations for Security Coverage, Model Inclusion and Classifications

Join us on LinkedIn

Click Here

Inside This Issue:

- Upcoming and Recent Events and Webinars
- Staff Speaking Engagements
- Attilio Meucci's Advanced Risk and Portfolio Management Bootcamp

Upcoming and Recent Events

2014 Northfield Annual Research Conference

The Stowe Mountain Lodge • Stowe, Vermont • October 5-8, 2014

Stowe Mountain Lodge

We are pleased to announce our 27th annual research conference at the Stowe Mountain Lodge, in Stowe, Vermont. The conference will officially start on Sunday, October 5th and end on Wednesday, October 8th.

For over a century Stowe has attracted the world's most discerning travelers, who have cherished the beauty of the area's covered bridges, sparkling lakes, rambling woodlands and majestic peaks. The Stowe Mountain Lodge has set a new standard of luxury and aesthetics for the East Coast's most magnificent resort.

The full seminar agenda and registration information will be posted to **www.northinfo.com/events.php** in the coming weeks as it becomes available.

Contact Kathy Prasad if you have any further questions, **kathy@northinfo.com**, 617.208.2020.

Northfield's Asset Owner Webinar Series Wrap-Up

Challenges Facing Asset Owners in Modeling Illiquid Assets as Part of an Overall Plan Portfolio

May 13th, 14th, 2014

Northfield's Asset Owner Webinar Series featured four presentations by Northfield's Dan diBartolomeo, Rick Gold, Emilian Belev, and a guest presentation by Mihail Garchev of PSP Investments. The topics addressed the challenges facing asset owners in modeling illiquid assets as part of the total plan portfolio.

The presentations included: "Factor Representation of Asset Allocation and Portfolio Risk Inclusive of Illiquid Investments," "Private Equity Real Estate Risk - The Good, The Bad, and The Ugly," "The Gorilla in the Room of Portfolio and Risk Management for Private Asset Classes: Business Cases for Institutional Real Estate and Infrastructure Portfolios" and "Unlisted Assets in the Context of Enterprise Risk Management."

The presentation slides are available at **http://www.northinfo.com/research.php** under the client conference proceedings section. Contact your Northfield Sales Representative if you are interested in viewing the full presentation recordings of the events.

Northfield's Private Wealth Webinar Series Wrap-Up

How Common Practice Falls Short of Best Practices

April 8th, 9th, and 10th, 2014

Northfield's Private Wealth Webinar Series featured three presentations by Northfield President Dan diBartolomeo that addressed critical topics in the management of private wealth. The material covered a range of best practices that benefit investors through superior after-tax returns, and also benefit asset management firms through greater operational efficiency. The presentations included: "Non-parametric Methods for Asset Allocation in Private Wealth," "Portfolio Management For Private Taxable Wealth: Basic Concepts and Operations" and "Generating Tax 'Alpha' for Private Wealth Households."

The presentation slides are available at **http://www.northinfo.com/research.php** under the client conference proceedings section. Contact your Northfield Sales Representative if you are interested in viewing the full presentation recordings of the events.

Attilio Meucci's Advanced Risk and Portfolio Management Bootcamp

August 11-16, 2014 • New York University • New York City

40 CE units CFA Institute, 40 CPE units GARP

The ARPM Bootcamp provides an in-depth understanding of buy-side modeling from the foundations to the latest advanced statistical and optimization techniques, in nine intense, heavily quantitative hours each day, with theory, live simulations, review sessions and exercises.

Topics include portfolio construction, factor modeling, copulas, liquidity, risk modeling, and much more.

Visit **http://www.symmys.com/arpm-bootcamp** to register, and view the detailed program information. There is a discounted Northfield supporter rate available. A short video is also available: **http://www.youtube.com/watch?v=BUnrgjNxBWk.**

2014 Newport Annual Summer Seminar Wrap-Up

Tennis Hall of Fame • Newport, Rhode Island • June 6, 2014

Northfield's annual summer seminar took place at the International Tennis Hall of Fame, in Newport, RI on June 6th. The seminar presented recent research and technical advances to an audience of Northfield clients and friends.

Tennis Hall of Fame

The agenda consisted of six presentations including: "An Integrated Multi-Period Approach to Risk and Asset-Liability Management," "Incorporating Commodities into a Multi-Asset Class Risk Model." "Non-Parametric Methods for Asset Allocation in Private Wealth," "Portfolio Optimization with VaR, CVaR, Skew and Kurtosis," "Risk Decomposition of Investment Portfolios" and "Risk Management for Public Pension Funds - Still Trying to Not Waste the Crisis."

As is customary, the seminar coincided with the USA Professional Championship of Court Tennis. Following the presentations, attendees viewed a Semi-Final Match between Camden Riviere and Steve Virgona. Court Tennis, or "real tennis" is the medieval sport that is the progenitor of all modern racquet sports. Riviere won the match and went on to win the finals. To learn more, visit the US Court Tennis Association site at **http://www.uscourttennis.org.** After tennis on Friday evening, everyone enjoyed a relaxing oceanfront dinner party at Oceancliff in Newport. The complete proceedings have been posted to our website at **http://www.northinfo.com/research.php.**

There is no charge for participation in any aspect of this event. We will accept any donation you might care to make on behalf of Pine Street Inn, Boston's primary homeless shelter. If you would like to make your donation online, please visit **http://www.pinestreetinn.org/donate/donate_now** or you can a make check payable to Pine Street Inn and mail to Kathy Prasad at Northfield. Should you have any questions please feel free to contact Chris Moy at 617-892-9172 or **chris.moy@pinestreetinn.org**.

Webinar Wrap-up: Introducing Illiquid Investments to Total Portfolio Management and ERM

April 30, 2014 • 12:00 P.M., E.D.T.

Northfield's Emilian Belev, CFA, ARPM, Head of Enterprise Risk Analytics, hosted a webinar where he discussed the rigorous risk analysis of illiquid investments like real estate and infrastructure alongside publicly traded investments in an ERM setting .

This was a Professional Risk Managers' International Association (PRMIA) event in coordination with the CFA Institute and qualified for 1 CPE Credit. Contact your Northfield Sales Representative if you are interested in viewing the full presentation recording of the event.

(Definition, Continued from page 1)

cox. The second would be Cremers, Kritzman and Page (2003) which did extensive simulations over a range of possible utility functions and found that the traditional mean-variance function is sufficient for almost all cases where the investor's utility function was not kinked (i.e. plausible bankruptcy).

Over any long period, the terminal wealth of an investor will come from three components. The original capital (zero return), the accumulated return on the capital (simple returns), and compounding of return on prior returns (geometric growth). For return magnitudes empirically observed, the third component will almost always dominate. To the extent that the magnitude of geometric returns is linearly related to variance, we would argue that it should be dominant risk concept. As such, we continue to prefer a decomposition of variance rather than standard deviation as the economically relevant measure of risk for *long term investors without liabilities at call.* Obviously this is somewhat of a semantic argument since if we ignore the degenerate root, there is a unique value of standard deviation for each value of variance and vice versa.

The other way to think about financial objectives is implicit in the way banks are run and regulated. You might describe this utility function as "maximize return subject to the constraint that the likelihood of insolvency in the current time period is less than P (something small like 1%)." Under certain restrictive assumptions (e.g. normal distribution and the ability to lever infinitely) this is equivalent to a utility function which maximizes the Sharpe Ratio. deGroot and Plantinga (2002) point out that such Sharpe Ratio maximization is consistent with plausible utility functions for only a small subset of investors (those willing to use significant leverage).

To the extent that banks have liabilities at call (depositors can withdraw their money at will), *it is natural to formulate the risk measure as dollar measure of solvency such as like VaR or CVar.* If we make random walk assumptions for asset values then VaR and CVar will be scalars of the standard deviation. *We would entirely concur that this makes sense for the kind of institution where day to day solvency is a primary objective.* For a sovereign wealth fund without any explicit liabilities there is no way to be insolvent so there would be little rationale for this kind of process, except perhaps to limit the reputational risk of staff and agents.

One can certainly argue that a defined benefit pension scheme has liabilities that are expressed as the present value of required future payments. However, those "liabilities" are really an actuarial representation of non-current, non-callable liabilities. In addition, the plan has the implicit asset of the payments guarantee from the sponsoring entity (a call option on future contributions). Our framing of the guarantee option was presented in **http://www.northinfo.com/Documents/556.pdf**. If the sponsoring entity happens to be sovereign government then plausible bankruptcy is ruled out, which again favors a mean-variance utility function per Cremers, et. al. Given the long term nature of the process, we do not see the merit of a solvency-driven risk measure being optimal for a defined benefit plan, *except in the special case that your plan takes the position that minimizing future contribution volatility is more important than minimizing the magnitude of future contributions.*

Risk Decomposition in Practice

Investment practitioners rely on a decomposition of portfolio risk into factors to guide investment decisions. While the total estimated risk of a given portfolio is usually unambiguous, the way in which risk is allocated to constituent factors (i.e. sources of risk) can be radically changed depending on the nature of the analytical and reporting process. The first issue is that for any model or framework for discussing risk, the statistical estimation process may vary even if defined with the same set of factors. A second source of ambiguity is that reporting conventions vary widely across vendors and systems. Another big issue is how the covariance among any pair of factors is allocated to the members of the pair. Finally, the inclusion or exclusion of basic portfolio constraints (e.g. portfolio weights should sum to 100%) will change the perception of risk sources. In addition, different metrics of risk (volatility, tracking error, VaR) may behave differently in consequence of the other issues.

Often the statistical estimation of a risk model will impact how risks are perceived to arise. All commercial factor models rely on a simple linear representation of asset (or portfolio) returns

$$R_t = \Sigma_{i=1 \text{ to } n} B_i F_{it} + e_t$$

R_t = the asset return in period t
B_i = the factor exposure to factor i
F_{it} = the return to factor i in period t
e_t = the return to factor i in period t

In times series models we observe the F values and statistically estimate the B values. In "fundamental" models we observe the B values and statistically estimate the F values. In blind factor (PCA) models we jointly estimate both. Let's assume we have two categories of factors called "red factors" and "blue factors." We could write such a model as to have G red factors and H blue factors

$$R_t = \Sigma_{i=1 \text{ to } g} B_i F_{it} + \Sigma_{i=g+1 \text{ to } g+h} B_i F_{it} + e_t$$

(Definition, Continued on page 5)

(Definition, Continued from page 4)

If we had a reason to do so, we could define the blue factors as being "net" of the influence of the red factors. For example, if inflation and interest rates were both factors in the model, we might choose to put inflation in the red group and redefine "nominal interest rates" as "interest rates net of the effect of inflation and other red factors (e.g. real interest rates). We accomplish this structuring of the decomposition by using a two-step estimation procedure:

$$R_t = \Sigma_{i=1 \text{ to } g} B_i F_{it} + z_t$$

$$z_t = \Sigma_{i=g+1 \text{ to } g+h} B_i F_{it} + e_t$$

z_t = the residual return at time t net of red factors only

Since we have defined and estimated the blue factors net of the red factors the risk decomposition *will naturally allocate more risk to the red factors and less to the blue factors.*

Why do "staged" model estimation? The usual answer is that you have a lot more data on some factors than others. Assume you have a universe of 1000 stocks broken into 50 industries. You will have 1000 data points for estimating the return to a factor like P/E or size, but only an average of 20 data points to estimate the return to a particular industry group. To the extent that some industries will have many participants (e.g. banks) and some will have very few (e.g. railroads) sample sizes on some membership factors may be very small. Another reason for staged estimations is that you have two or more factors that are highly correlated. Statistical estimation procedures often produce unstable results when independent variables are correlated. By defining one factor net of another correlated factor, we structurally remove their natural correlation.

Finally, we may prefer staged estimation for particular strategies where it makes sense. For example, there has been a long debate about whether countries or sectors are more important to global equity portfolios. The answer depends on whether you see the world as cap weighted or equal weighted. If you see the world in a capitalization weighted fashion, a large part of the world's big firms operate globally and compete against global competitors. The influence of country effects is usually small. If you see the world in an equal weighted way, there far more small local enterprises than there are big global enterprises. As such, you would perceive country effects as dominating.

The second source of ambiguity in risk reporting is the conventions of the various vendors of risk estimates. At Northfield, we always prefer to decompose expected variance of returns. Variances are *naturally additive* while standard deviations (or VaR segments) are not. Variances are also consistent with a long-term view of risk as decreasing compounded returns relative to arithmetic returns, as discussed in the introduction. It is typical to then decompose the variance into factor and asset specific components. For a model with N factors, we will have a factor covariance matrix of N*N elements. This matrix is square and symmetric about the diagonal. We use the conventional assumption that of the covariance between any two factors, we will credit half of the covariance to each factor. There is an algebraic (not but necessarily economic) rationale for this. We then create a **row subtotal for each factor and report it .**

Depending on which Northfield model is being used, factor exposures may or may not be put on a common scale. For non-scaled exposures the sign of an active factor exposure is relevant for both benchmark relative and absolute risk. Scaled exposures (e.g. Z scores) are often difficult to interpret in terms of absolute risk or VaR since the exposures are relative to the central tendency of the asset universe. Under our reporting, the magnitude of an active "bet" must be judged from factor contribution in **variance units.** Security specific risk is summed across positions and presented as a single value. For our multi-asset class "EE" model, the relationships of multiple securities from the same issuer (e.g. Bank of America stock and a Merrill Lynch bond) are accounted for properly.

Another popular risk vendor decomposes variance into both factor and specific components but with a very different scheme of creating sub-totals. Each factor variance contribution is reported separately. *All factor covariance terms (the off diagonal elements) are added up to a single sum* and are reported separately. This obfuscates the joint effect of two or more factors that are correlated. All factor exposures are scaled (Z score or percentage). Benchmark relative active factor exposures can be easily observed and are often incorrectly interpreted as being the relative measure of bet size in either **standard deviation** or variance units **as the factor variances are not uniform.** Neither the sign nor the magnitude of factor exposures may be easily interpreted in absolute risk terms (i.e. a riskless asset like cash doesn't have a market cap or PE).

Yet another vendor decomposes variance into factor and specific risk. Statistically defined factors are assumed orthogonal so no off-diagonal covariance terms exist. These statistical factors must be mapped onto real world factors for economic interpretation. All factors are defined to have the same unit volatility. Factor exposures are rescaled to reflect the relative risk of each factor in **standard deviation units**. A unique feature of statistical models is that the signs on factor exposures are arbitrary [1*1=1= (-1*-1)]. We can define continuous factor exposures in an arbitrary fashion but we find it very unintuitive to do this with something like industry memberships. Your factor exposure to the "short oil industry factor" is the negative of your factor exposure to the "oil industry factor."

(Definition, Continued on page 6)

(Definition, Continued from page 5)

As previously noted, we assume that across any pair of factors (or assets), half the covariance is allocated to one factor and half is allocated to the other. Let's do a simple two asset example in absolute variance units.

Stocks - 20% Volatility, 60% weight

Bonds - 5% Volatility, 40% weight

Assume 20% correlation between Stocks and Bonds

The variance of the portfolio is:

$V_p = 20^2 * .6^2 + 5^2 * .4^2 + 2 * 20 * 5 * .6 * .4 * .2 = 157.6$

$V_p = 400 * .36 + 25 * .16 + 2 * (20 * 5 * .6 * .4 * .2) = 157.6$

The conventional thing to do is to allocate half the covariance of a pair (in red) to one asset (or factor) and the other half to the other asset (or factor). Rearranging can get:

$V_p = [144 + (20 * 5 * .6 * .4 * .2)] + [4 + (20 * 5 * .6 * .4 * .2)]$

$V_p = 148.8 + 8.8 + 0$

Alternatively we can get

$V_p = 144 + 4 + 9.6$

Allocated to stocks, Allocated to bonds, Allocated to covariance

In the first decomposition, the total covariance was 9.6, of which 4.8 was added to the first term (stocks) and 4.8 has been added to the second term (bonds). It's easy and algebraically simple but not necessarily indicative of any underlying economics. Our question is whether it is economically realistic to do this "half and half" split given that the variance contribution of the stocks alone (much more volatile and bigger weight) is 36 times as big as the variance contribution of the bonds alone (much less volatile and smaller weight)?

One could easily argue for some form of proportional structure like this:

$V_p = [144 + (144/148 * 9.6)] + [4 + (4/148 * 9.6)]$

$V_p = [144 + 9.34] + [4 + .26] = 153.34 + 4.26$

In the first formulation, the amount of variance allocated to bonds (e.g. factor 2) is more than double what it is in the second and third formulations. If I were the bond portfolio manager, I'd like the second formula a lot more than the first. The conventional approach (half and half) is one end of a spectrum. The Vendor 2 approach of keeping covariance completely separate is the other end of that spectrum. There is a lot of area in between where some kind of proportional allocation makes the most sense. We have proposed other proportionality procedures (e.g. allocating based on the relative absolute value of the product of weight and volatility) but *results are still somewhat arbitrary.*

Many systems (e.g. Vendor 4) try to decompose risk (either variance or standard deviation) by the *security position.* This is often perceived to be intuitive for VaR calculations at banks (e.g. how much risk comes from each loan) because the *alternative is to not make the loan*, and we are measuring risk of loss in dollar amounts. For asset management, the amount of capital (AUM) to be put at risk is fixed (i.e. institutional clients don't pay asset managers to hide money under their mattress). Implicitly our risk calculations must decide whether we do or don't want to enforce the requirement that asset weights sum to 100%. As such, any algebraic decomposition of risk by position requires (either explicitly or implicitly) the definition of a "contra-asset" which defines where the proceeds of closing out a position will be deployed. For analysis of absolute risk or VaR, the usual algebra implicitly defines the contra-asset as riskless cash. For benchmark relative decomposition of incremental tracking variance by position, you could define the contra-asset as either cash, an ETF for the benchmark or void (reweight the remaining portfolio positions to again add to 100%)

Choosing cash is often done but often causes confusion in asset management because it's easy to have positions that are big risk contributors in absolute terms (e.g. high beta) but diversifiers on a benchmark relative basis (high beta is diversifying if the rest of the portfolio is low beta relative to benchmark. Implicitly going to cash usually increases perceived tracking error while reducing absolute risk. Using a benchmark ETF is messy. If the position X you are selling out is a member of the benchmark, selling out position X and replacing it with the benchmark ETF implicitly buys back some of the stock X risk exposure that you just thought you got rid of. The problem becomes mathematically recursive.

You can act like a bank calculating it's $ VaR. Reweighting the remaining positions for asset management creates confusion because the new portfolio weights you will end up with after selling out position X, or position Y will be different. As such, you can't directly compare the incremental volatility or variance risk changes across positions, which defeats the purpose. This approach often works well for VaR because removing the successive positions and reweighting reflects that economic value of your portfolio has declined by the correct increment. However, VaR is an *incoherent measure*. You can show that it leads to clearly wrong conclusions about risk for some problems. Some choices of "contra-asset" allow closed form allocation of volatility (standard deviation) and VaR and others do not.

From an analytical perspective, the only risk quantities that are exactly known are the *marginal variances* (MV) of a factor or position. The problem is that the marginal variance has to be obtained from basic calculus (which non-math people don't like). The MV values are legitimate only for infinitesimally small changes in position size which non-

(Definition, Continued on page 7)

(Definition, Continued from page 6)
quant people see as unintuitive and non-actionable. However, the marginal variances, not the incremental contributions by position are what matter in defining optimality (Kuhn Tucker conditions). For most assets, when you consider trading a position you can close out any part of it, not just "all or none." The exception to this assertion would be something very illiquid like real estate (i.e. tough to sell half a building).

The algebraic issues get even worse when you try to decompose by position in standard deviation units. This is widely done to make it easy to do parametric VaR by position. Since standard deviations are not naturally additive you have to use some kind of clever algebra to allocate SD risks by position to add to the SD total. Some of the proposed schemes distort the economics less than others but *there is no exact solution for some definitions of the contra asset.* Many systems do the variance contributions and then divide everything by the standard deviation *as a scalar constant,* which creates percentage allocations of standard deviation that are identical to the percentage allocations by variance. Some vendors try to decompose by position, and then by factor within position. This produces lots of numbers that add up to the portfolio volatility but are very hard to use to actually make portfolio decisions.

A clever solution is presented in Menchero (2010) which he refers to as "X, Sigma, Rho." The goal is to minimize distortion when trying to allocate volatility (e.g. tracking error) to additive increments per position that sum to the correct total. X is the exposure of the portfolio to a security, sigma is the volatility of the security and rho is the correlation of the security to the existing portfolio (not the benchmark). You can think of the risk of each position as having two components, one that is 100% correlated with the risk of the portfolio, like a parallel line segment and a second component that is correlated at some value (obviously bounded between -1 and +1). The decomposition can be geometrically interpreted as if the correlation between the security and the portfolio was an angle of which we are interested in the sine and cosine. While this works geometrically, *the contra-asset problem remains.* Rho is the correlation between a particular stock X and the portfolio in which an active position in stock X is already present. If I take stock X out of the portfolio, something else has to go there. The portfolio is now different, all the correlations will be different and so will *all the lengths of the parallel line segments.* The trigonometric process to make the incremental position risk contributions additive does not make them independent.

Conclusions

For most long term investors, the decomposition of risk should occur in variance units rather than volatility related units. For financial intermediaries *with liabilities at call* where solvency is a predominant issue, the opposite is true.

How much risk we allocate to a given factor is heavily influenced by the estimation process of the model. The robustness of statistical estimates can often be improved by staged estimations, but at the cost of more complex interpretation. Risk service vendors report the decomposition of risk differently. Many of the reporting procedures follow an algebraic rather than economic reasoning. Much of the ambiguity relates to how the covariance terms are allocated to the involved factors.

When dealing with "incremental risk contributions by position" we will be either implicitly or explicitly dealing with the existence of the contra-asset. Only some of the possible definitions of the contra-asset have simple algebraic structures.

Staff Speaking Engagements

On June 11th, Northfield President Dan diBartolomeo was at the Joint Meeting of CFA, PRMIA and QWAFAFEW in Stamford, CT, where he presented "Non-parametric Asset Allocation for High Net Worth Individuals."

Dan will be presenting at the Credit Suisse Quant Conference in New York City on June 25th. The topic will be "Fat Tails, Tall Tales, Puppy Dog Tails."

Dan will present "Analytical Processes for Evaluation of Hedge Funds," at the CFA Brazil meeting in Sao Paulo on August 7th.

Dan will be in Oxford, UK on September 8-10 for the London Quant Group Annual Conference where he will be presenting "Are Long Term Returns Predictable in Efficient Markets?"

Dan will be presenting "Risk Monitoring for Plan Sponsors" at the Milliman Client Conference in Lake Tahoe, CA, on October 2nd.

Northfield's Louis Scott presented "Microstructure, Trading Costs and Optimal Algorithmic Execution at the Optirisk Workshop on Market Microstructure in London on June 16th.

Louis presented "Sentiment, News and Volatility in Stock Returns" at the Optirisk Conference on Sentiment Analysis in Finance on June 19th.

Northfield Asia's Nick Wade will be presenting on Risk Model Design at the RMIT Finance Day in Melbourne on August 22nd.

Technical Support Tip: Analytical Considerations for Security Coverage, Model Inclusion, and Classifications

By Steve Dyer

An important step in the optimization process is ensuring that your data is complete and accurately represents both your actual holdings as well as your beliefs about your holdings. This Tech Tip will discuss the most important data considerations clients need to understand to ensure accurate optimization results.

Missing Securities

It is often the case that securities, even very common securities, are not included in the model, or that they are intentionally excluded from coverage. Northfield deliberately excludes a security from the data set when its properties violate assumptions for a linear model to preserve the integrity of the model. As a matter of course, there are numerous criteria by which we filter securities from our routine coverage universe. There are minimum limitations on both market capitalization and trading volume. A firm must have at least $100 million in market capitalization to be included. There must be non-zero trading volume on the last trading day of the month to verify that the listed closing price is not too stale. We also exclude securities where trading has been suspended on the principal exchange (this is usually indicative of an allegation of accounting fraud and therefore unreliable underlying data). For new securities that otherwise meet requirements, they are included in all models under the normal process after two month-end price observations are available, with the exception of the Fundamental model. Since we can directly measure exposures to fundamental factors, we can include new securities in the Fundamental model immediately in the month following their issue. If clients require coverage of new securities sooner than this, we encourage them to submit a request to the support desk and we will generate an exposure record using available data.

Northfield will also routinely exclude securities for caution's sake, for example if a company has extraordinary circumstances that would indicate its return history does not reflect its current risk. This does not mean we cannot include the security in the analysis of a portfolio by a specific client, but rather we do not want clients using questionable data without being aware of the issues at hand. Companies that are being excluded typically are experiencing uncommon conditions and need to be handled on a case-by-case basis. In these situations, we encourage clients to contact support (support@northinfo.com) to discuss the excluded security.

ADRs and GDRs

To the extent that our risk models involve classifying firms into regional or national memberships (e.g. USA or Europe) there is an interesting problem posed by the existence of ADRs and GDRs. In our models, ADRs are classified as dollar-denominated American securities; GDRs are similarly treated as securities denominated in the local currency and as a member of the local market. Most risk systems simply map ADRs and GDRs back to the underlying local share, but we believe that this is an easy but *inaccurate* solution in many cases. There are many cases where mapping to the local underlying security is the obvious choice – for example, an investor owning the ADR for the Japanese telecom firm Nippon Telephone would want to treat it as a Japanese security, as most shareholders are in Japan, most of the trading volume is in Japan, and most of the business revenue is in Japan. On the other hand, with multinational companies, it is not always appropriate to map to the local underlying security. For example, Royal Dutch Shell's primary listing is in London, but most of its revenue is in the United States, and oil prices are denominated in dollars, so the ADR might more appropriately represent its actual market and currency risk. There are even more ambiguous cases, such as mining companies, that might operate mines in Africa, are headquartered in Europe, but trade mostly in New York. Because there is considerable disagreement about how to handle these cases, we leave it to the responsibility of the user to make these judgment calls and make decisions appropriate to their investment beliefs.

Northfield systems allow the use of the composite asset function to permanently map each ADR to the local shares, or the local shares to ADRs as a user thinks appropriate. The **screenshot below** shows how to remap the primary listing of Royal Dutch Shell to the ADR as in the example above (Please note that if a user creates a composite asset with the same ID as already exists in the database, the

(Tech Tip, Continued on page 9)

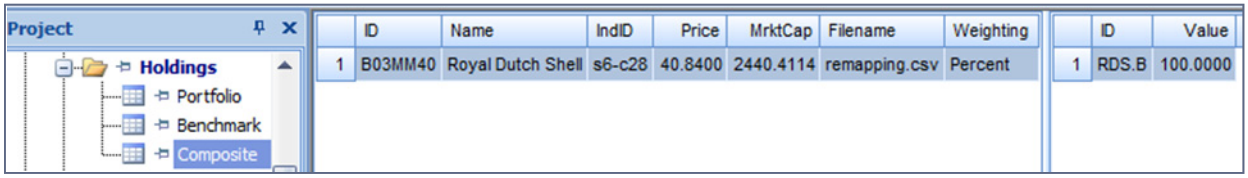

(Tech Tip, Continued from page 8)

Optimizer will recognize the composite asset and ignore the data record. The underlying assumption is that the composite asset was created by an intentional, manual process, so it must be what the user wants).

This is a critical issue when mixing ADRs and local shares in the same portfolio analysis, as often arises when a benchmark index is denominated in local shares but a US investor will buy ADRs for legal convenience or by mandate. This is because the model assumes that all specific risk is uncorrelated, so in addition to different exposures to factors, the system does not recognize the specific risk of an ADR and the underlying security to be related, resulting in an upwardly biased risk estimate.

Multiple Share Classes

As an extension from the ADR discussion, the default model behavior similarly treats multiple share classes (e.g. Google Class A vs. Google Class C, common vs. preferred) as different securities with uncorrelated specific risks. Again, this is intentional, as the different share classes' specific risk values are roughly the same only from a bankruptcy risk perspective. Our estimate of specific risk includes several other factors including the presence of kurtosis in the return distribution (fat tails), evidence of serial correlation in the returns (illiquidity) and especially in the case of Google's new share class, the extent to which any statistical measure can be relied upon using limited data of only a few months of returns. It is our practice to always take a conservative view of any situation of particular uncertainty or reduced data availability and to overestimate the risk of a security rather than underestimate it. However, if a client wishes to override the way different share classes are handled, they can map the share classes to each other in the same manner as described above.

Secondary Exchanges and Cross-Listings

One step further in complexity from the Nippon Telephone/Royal Dutch Shell example are situations where firms have listings and large trading volumes in markets where the firm does not have primary business operations. These cases generate natural disagreements among investors and clients related to time horizon – long-term investors think of risk in terms of risk faced by the underlying business, while short-term investors think in terms of how the investors that are involved in the share price setting process are behaving. A recent instance that illustrates this philosophical divide was a Russian IT firm that was legally incorporated in the USA, and where essentially all the trading volume occurred on the US NASDAQ. In this case, the long-term investor could argue that the operating risks of the business are in Russia, so the national designation and currency risks are associated with Russia. The short-term investor could equally argue that the security returns were arising from share prices changes were arising largely from trading by US investors, not Russian investors, so the market risks of holding the shares relate to the US market.

This is an instance where it would be incumbent upon the investor to make a choice about his or her investment beliefs about the firm in question. In most cases, mapping the firm using a composite asset to something sensible is a sufficient, permanent solution and works across all models, but the choice of what mapping to choose is dependent on the investing philosophy of the investor.

Technical Support in Boston can be reached at: 617.208.2080, **support@northinfo.com.** European clients can contact: **support-europe@northinfo.com** or call +44-(0)-20-3714-4130. In Asia, call +81(0)3 5403 4655 or +61(0)2 9238 4284 or **support-asia@northinfo.com.**

If you have any suggestions of what you would like to see covered in upcoming issues, please e-mail your ideas to general@northinfo.com

For a complete index of all former Northfield News articles, visit http://www.northinfo.com/documents/314.pdf

Boston Office
77 North Washington Street, 9th Floor
Boston, MA 02114
Phone: 617.451.2222
Fax: 617.451.2122
Sales: 617.208.2050
Tech Support: 617.208.2080

London Office
2 - 6 Boundary Row
London, SE1 8HP
Phone: +44 (0) 20 3714 4130

Tokyo Office
Shiroyama Trust Tower
4-3-1 Toranomon
Minato-ku
Tokyo 105-6027
Phone: +81 (0)3 5403 4655
Fax: +81 (0)3 5403 4646

NORTHFIELD *Northfield News* is a publication of Northfield Information Services, Inc., 77 North Washington Street, 9th Fl., Boston, MA 02114. If you have any questions or comments regarding the content of this newsletter, please call us, or e-mail us at staff@northinfo.com, or visit our home page at http://www.northinfo.com