December 2010

Northfield News

A Newsletter for the Friends and Clients of Northfield Information Services

Portfolio-centric Algorithmic Execution of Equity Trades

By Dan diBartolomeo

Unknown to most clients, Northfield has been involved in the creation of trade execution algorithms since 2004. Algorithmic execution of buy-side orders has steadily gained an ever larger share of trading volume in most equity markets around the world. At the same time, the provision of liquidity from high-frequency trading operations has expanded even faster. Taken together, these two developments demand ever-increasing sophistication in execution algorithms. To that end, Northfield has created two new mechanisms for algorithmic execution. The first is a "pre-processor" that aligns the parameters of the execution algorithm to the composition and strategy of the underlying portfolio. The second is an enhanced version of our existing "trade scheduling" algorithm that breaks large orders ("parent order") into a series of smaller trades ("child order") to be executed over time. Our approach to optimal execution is based on a multi-period mean variance optimization in discrete time. The first implementation of these two techniques is in the new ITOS service being offered by Algorithmic Trading Management (ATM), a leading provider of order execution technology for equity traders.

We assert that the first key to successful use of execution algorithms by asset managers is to have effective communication between portfolio managers and traders so that the

Portfolio-centric, Continued on Page 4

Intra-Horizon Risk

By Nick Wade

The concepts of tracking error and value-at-risk (VaR) are widely adopted as risk management measures within the investment community. However, they are not free of criticism. For example, both tracking error and VaR have been criticized widely in the literature for the lack of information they provide about tail risk. One of the key assumptions underlying both tracking error and VaR is that security or portfolio returns are Normally distributed. The distributions of return observed in markets – particularly at higher frequencies such as daily or intra-day – are non-Normal; they exhibit fat tails and often skew. The use of a measure that relies upon the Normal distribution therefore underestimates the size of the possible loss to the extent that the actual distribution deviates from Normal. Various measures have been proposed to resolve that issue.[1]

The concept of "tracking error" came out of index fund management as a way of quantifying the extent to which a particular fund deviated from an index. In active management, as the differences from a particular benchmark become more and more marked, so the relevance and usefulness of tracking error in its unmodified form as a risk measure decrease. For example, see the recent presentations by diBartolomeo on "Strategy Risk."[2]

In addition, reliance on the Normal distribution ignores the effects of jumps in returns or

Intra Horizon, Continued on Page 8

Special Points of Interest:

Join us on LinkedIn

Click Here

Inside This Issue:

Recent and Upcoming Events

Northfield Workshop: Lies and Performance Attribution

January 25, 2011 • 11:00 A.M., E.S.T.

Northfield's Steve Gaudette will be hosting a webinar presentation which will discuss approaches for uncovering the truth behind performance attribution, on Tuesday, January 25, 2011 at 11:00 a.m., Eastern Standard Time.

There is no charge to participate. To register, and view the detailed abstract for the presentation visit https://northinfoevents.webex.com.

Northfield Webinar: Key Elements of Risk Control for Asset Managers

March 8, 2011 • 11:00 A.M., E.S.T.

Northfield's Richard Pearce will be hosting a webinar presentation which will discuss risk management, on Tuesday, March 8, 2011 at 11:00 a.m., Eastern Standard Time.

There is no charge to participate. Registration will be opened in February. To register and view the detailed abstract for the presentation visit https://northinfoevents.webex.com.

Northfield Webinar: Redefining Private Equity Real Estate Risk

December 7, 2010

Northfield's Rick Gold and Emilian Belev hosted a webinar presentation on Redefining Private Equity Real Estate Risk , on December 7th. The presentation slides have been posted to http://www.northinfo.com/documents/407.pdf

The event recording can be viewed at:
https://northinfoevents.webex.com/northinfoevents/lsr.php?AT=pb&SP=EC&rID=3015022&rKey=cb9fd812072e071e

---

Northfield Webinar: The Central Paradox of Active Management

October 19, 2010

Northfield President Dan diBartolomeo hosted a webinar presentation on The Central Paradox of Active Management, on October 19th. The presentation slides have been posted to http://www.northinfo.com/documents/399.pdf

The event recording can be viewed at:
https://northinfoevents.webex.com/northinfoevents/lsr.php?AT=pb&SP=EC&rID=2828192&rKey=5770e620a9293584

Northfield Annual Holiday Party Wrap-up

Boston • December 15, 2010

Clients and friends joined Northfield for our annual Holiday Party on the evening of December 15th.

Complimentary cocktails and and Hors d'oeuvres were served. Live entertainment was provided by an a local acapella quartet from the New England Conservatory of Music.

This is an annual event which is open to all. Registration goes live in November, so check the Northfield website next year if you missed this year's party.

Northfield European Seminar Wrap-up

Le Méridien, Piccadilly, London • November 9, 2010

The Northfield 2010 European Investment Seminar was held in London at the Le Méridien Hotel on November 9, 2010. The purpose of the seminar was to highlight recent advances in analytical techniques for the investment industry to our growing number of European clients and prospects.

The presenters included Northfield's Dan diBartolomeo, Rick Gold, Anish Shah and Nick Wade. Guest speaker Steve Satchell of Trinity College, Cambridge University, also gave a presentation. The topics included; "An Introduction to Independent Components Analysis (ICA)," "Equity Risk, Credit Risk, Default Correlation, and Corporate Sustainability," "Intra-Horizon Risk," "Real Estate: A review of real estate's contribution to portfolio risk and return in the new world financial (dis)order," "The Discretionary Wealth Hypothesis in an Arbitrage-Free Term Structure to Asset-Liability Management" and "The Sensitivity of Beta to the Time Horizon when Log Prices follow an Ornstein-Uhlenbeck Process."

The seminar concluded with a well deserved post seminar reception. There was no cost to attend, however, donations to the Prince's Trust were strongly encouraged. The Prince's trust is a very worthwhile organization that makes a huge positive difference to the lives of many thousands of young people. Visit **http://www.princes-trust.org.uk** to learn more.

The seminar proceedings have been posted to **http://www.northinfo.com/papersearch.cfm.**

Asia Seminars Wrap-up

Hong Kong, Singapore, Sydney and Tokyo ● October 2010

Northfield hosted our annual Asia Seminar Series with four highly successful events in Hong Kong, Singapore, Sydney and Tokyo. The seminars showcased our research on key topics in investment and risk management to our growing family of Australian and Far Eastern clients and prospects and broadened awareness of the range and depth of Northfield products, services, and research.

The presentations were given by Northfield's Dan diBartolomeo, Nick Wade, James Williams. and guest speaker Lloyd Kurtz of Nelson Capital. Topics included: "Extending Factor Models of Equity Risk to Credit Risk and Default Correlation," "Improving Portfolio Construction through Adjustment for Parameter Estimation Error," "Incorporating Strategy Risk of Active Managers into Portfolio Risk and Optimization," "Incorporation of Liquidity Risks into Equity Portfolio Risk Estimates," "Intra-Horizon Risk," "Personal Asset/Liability Management: Using the Discretionary Wealth Hypothesis within an Equilibrium Term Structure" and "The Stakes Go Up in Social Investing: New Evidence, New Controversies."

Complete seminar proceedings will be posted at **http://www.northinfo.com/papersearch.cfm.**

Asset Allocation and ART Product Seminar Wrap-up

Tokyo ● December 14, 2010

The recent Northfield Asia Product Seminar on Asset Allocation was held in Tokyo Japan on December 14th with approximately 50 attendees from asset managers, security brokers, pension consultants etc. This year's seminar focused on Asset Allocation with the Northfield Allocation Research Toolkit (ART) service, and covered investment suitability, manager selection and monitoring, style analysis, and optimization, introducing the analytical hierarchy process, CUSUM process-control, mean-variance and mean-semi-variance optimization incorporating Bayesian adjustments for estimation error, and style analysis importantly including confidence intervals around style weights. Yasuhiko Nakase of Northfield Asia led the discussion. Reflecting the high level of interest in the subject matter. the Q&A session was exceptionally active with serious questions.

(Portfolio-centric, continued from page 1)

parameterization of the trading algorithm can be tailored to the strategies and objectives of the underlying portfolio. For example, an alpha strategy predicated on response to earnings surprises may require very urgent trading, while a much more passive approach to trading would be appropriate for a portfolio that is making a trade based on a long term view of a particular stock's growth potential. Similarly, trading algorithms should take into account whether the trades come from a very aggressive or risk-averse portfolio and whether the portfolio strategy has any explicit view on serial correlation of security returns (i.e. a value or momentum tilt). The unfortunate truth is that within many asset management organizations there is minimal communication between portfolio managers and traders, often limited to the form of a basic order ticket.

The lack of communication between portfolio managers and traders is most often manifested by trading that is too aggressive, *with traders paying unnecessarily high market impact costs in order to execute trades faster than the economically optimal pace.* One study that documented this effect was done by Instinet and presented at our 2007 Newport seminar, **http://www.northinfo.com/documents/246.pdf**. This study tracked more than 20,000 program trades done by Instinet clients. Some trades were done using a "volume weighted average price" (VWAP) algorithm, others were done using a "percent of volume" (POV) procedure, and others were done with a previous version of the Northfield trade scheduling algorithm that had been licensed by Instinet. The study found that trades done using the VWAP and POV procedures were consistently more expensive. Of course, more urgent trading could be justified if the alpha the manager expected to gain was sufficient. However, the added costs of the conventional algorithms implied an expected alpha of more than twenty basis points per day, or about 50% per year. Similarly, one could justify very aggressive trading as an outcome of high risk aversion. However, it was found that to justify the higher cost of more rapid trading an investor would have to be so risk averse that they would never be willing to invest in equities in the first place. The dominance of the Northfield multi-period approach was statistically significant for parent orders (trades) large than 6% of average daily volume in a given stock.

To address the sub-optimal parameterization of trading algorithms, Northfield has created a "pre-processor" tool that uses our risk models to analyze the portfolio from which the orders arise. The analysis performed provides three scalar values and one vector as inputs for trading algorithms. No other information on the underlying portfolio is disclosed into the trading process. In the ITOS implementation, both the pre-processor and the trade schedule computation engine sit at the buy-side client site so as to remove any potential for inadvertent leakage of position information. **(see fig. 1).**

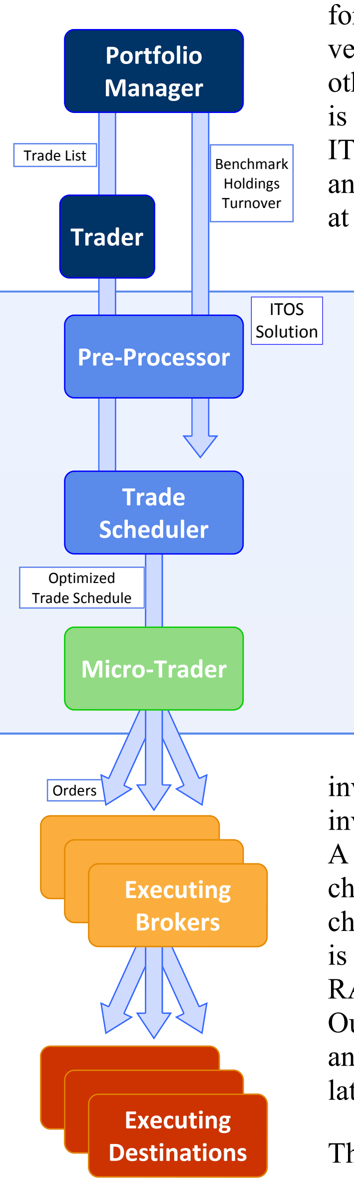

Figure 1

The first item analyzed is how aggressive or conservative is the underlying portfolio. In Northfield terminology, this property is quantified as the Risk Acceptance Parameter (RAP). If you have an explicit alpha forecast for the portfolio we can just calculate the slope of the mean variance efficient frontier as you lever the portfolio slightly. For example, we can see how expected alpha and tracking variance change as you go from being 99.9% invested with .1% investment in the benchmark to being 100.1% invested with -.1% invested in the benchmark. A finite difference estimate of RAP is just the change in tracking variance divided by the change in alpha. If no explicit alpha forecast is available, we can make an inference of RAP from only the portfolio tracking error. Our approach is adopted from Wilcox (2003), and results in a simple rule of thumb that relates RAP to the portfolio tracking error.

The next output from the pre-processor is the vector of "implied alphas". Once we know the risk tolerance of a given portfolio, a set of alpha forecasts must exist that will make the portfolio mean-variance optimal. These are the "implied returns" for the portfolio, as first described in Sharpe (1974). For long-only portfolios an adjustment is typically required to the basic implied return calculation. Our technique for this adjustment is described in diBartolomeo (2008).

(Portfolio-centric, Continued on page 5)

(Portfolio-centric, continued from page 4)

Once we have the implied alphas at the portfolio level, we can apply an alpha decay rate and work backward to a vector of "short horizon" alpha estimates that would be relevant during the period of the trade execution. Unless supplied by the portfolio manager, we assume an alpha decay rate consistent with Grinold and Stuckelman (1993), where they derive the approximate relation that "half the turnover will get you three-fourths of the value added." With an expected level of annual turnover and some algebra you can turn this into an alpha decay rate and estimate "per day" alpha decay on your implied alphas.

The final parameter obtained from the pre-processor is a measure of the implicit assumptions of the underlying portfolio with respect to serial correlation in security returns. Relevant Northfield risk factors are used to characterize the portfolio. If a portfolio is momentum oriented, the manager is expecting that stocks that have gone up will keep going up and stocks that have gone down will keep going down (positive serial correlation). On the other hand, a value oriented investor is generally assuming that stocks that have gone down will come back up and those that have risen will subsequently decline (negative serial correlation). For more details, see diBartolomeo (2003).

Depending on the expectations of serial properties, most trading algorithms offered by brokerage firms will either speed up or slow down their rate of trading in response to changes in security prices relative to the price that existed at the start of the execution process. Traders often refer to the changing rates of orders executed per unit time as a "trajectory."

The full output of the pre-processor can then be fed into the recently updated Northfield trade scheduling algorithm. *You can think of a list of undone trades as a long/short portfolio that you need to liquidate to an all cash portfolio. You are long things you do have and don't want (sell orders). You are short things you do want and don't have (buy orders).* We have two motivations to trade quickly. The first is opportunity cost or short term alpha. If we are an active manager buying a stock, we are presumably doing so because we believe the stock will go up, and we wish to buy before it goes up, not after. If we are selling, we believe the stock price will fall, and we wish to transact before, not after the decline. The second reason to trade quickly is risk. The longer the trades remain unexecuted, the longer our underlying portfolio is different from what we want to hold, and we bear uncertainty around the relative performance of the portfolio we actually hold at each moment in time compared to our desired (or target) portfolio.

On the other hand, we have a reason to trade slowly. The faster we demand liquidity from the market, the greater our market impact on prices will be. To the extent we cannot do all of our trades in a single execution, our buying will drive up prices making our subsequent purchases more expensive. If we are selling, we will drive down prices affording us lower prices for our subsequent sales. By framing the execution of undone orders in a portfolio context, we can use an extension of our existing portfolio optimization methods to optimally balance short term alpha, risk and market impact costs in a fashion consistent with the composition and strategy of the underlying portfolio.

Our algorithm is a multi-period mean variance optimization in discrete time. We define N periods per day such that each successive period has an expectation of (100/N) percent of the daily volume. Typically, the number of minutes in a period is smaller at the beginning and end of the day (busier trading) and larger during the middle of the day (slower trading). An easy way to visualize the output trading schedule is a spreadsheet where orders are rows and time periods are columns. The sum of any row is the entire order for that security. The whole schedule can be recalculated in a few seconds if market conditions or other inputs change.

A mathematical representation of the problem looks like this:

$$U = E\,[\Sigma_{t=1 \text{ to } N}\,(\Sigma_{i=1 \text{ to } m}\, w_{it}\,\alpha_{it} - (1/RAP)\,\Sigma_{i=1 \text{ to } m}\,\Sigma_{j=1 \text{ to } m}\, w_{it}\, w_{it}\,\sigma_i\sigma_j\rho_{ij}) - \Sigma_{i=1 \text{ to } m}\,(\text{abs}\,(w_{it}-w_{it-1}))\,\kappa_{it}]$$

U=expected value of the objective function that we want to maximize

N=number of time periods in the schedule

w_{it}=the weight of security i in the "undone" portfolio at the start of period t

α_{it}= the short term alpha associated with security i during period t

ρ_{ij}=the correlation of security i and security j

σ_i =the standard deviation of returns to security i

κ_{it} =amortized percentage trading cost for security i in period t

E =the expectations operator

The optimization problem is to obtain the matrix of security weights over the set of both securities and time periods (w_{it}) such that the objective is maximized, subject to the constraint that at the completion of the schedule in N periods the starting weights for the hypothetical next period (w_{n+1}) for all non-cash securities must be zero.

(Portfolio-centric, continued on page 6)

(Portfolio-centric, continued from page 5)

It should be noted that the expected value for U will almost always be negative. In this case, our alpha values are reversed from a conventional optimization, as a long position in this "undone" portfolio represents a sell order which is apt to be associated with a negative alpha, and a short portfolio position represents a buy order which is apt to be associated with a positive alpha. The other two terms are both negative by construction.

The computational challenge of this problem is two fold. The first issue is the multi-period nature of the problem which induces path dependence in the solution. If you sold 100 shares of stock XYZ in the first period, you don't have those shares any more to sell in the second period. The second challenge is that the matrix of trading cost values, (κ_{it}) is a complex function of the amount each stock that is to be traded in a given period. It is also path dependent because we assume that some portion of market impact is permanent and accumulates across multiple periods. Finally, we also take into account that the market impact of trading a particular security will affect the price of other securities. For example, buying a large amount of Ford stock will push up the price. To the extent that investors base their valuation of GM and Toyota at partially on the price of Ford, the expected prices of GM and Toyota are also changed. This means that a combined order to "buy Ford, buy GM" will have a different expected cost than a combined order to "buy Ford, sell GM" executed in the same period. The conventional computational approach to addressing the path dependence would be to use dynamic programming, as in Bertsimas and Lo (1998), which is often very slow to compute for large problems. We have developed an alternative approach that allows us to solve a large problem (e.g. 200 parent trades done over 20 periods) in a few seconds.

We believe despite the complexity, our approach to trade scheduling has very desirable properties compared to popular algorithms. Most importantly, it explicitly controls risk of the "portfolio" undone trades at each moment in time during the execution process. Secondly, the schedule can be constrained such that the cash balance is always positive (never overdrawn). Finally, the distribution of expected total costs is naturally symmetric, but can be skewed to either side as needed to be consistent with portfolio strategy.

Many brokerage firm algorithms are just ways of adjusting trading speed in response to real time price changes. The net cost of an execution relative to the arrival price is typically referred to as "implementation shortfall" as first described in Perold (1988). If the current price is favorable relative to the price at the start (arrival), you speed up relative to some originally prescribed trajectory of transactions. If the current price is unfavorable relative to the arrival price, you slow down relative to original trajectory. The problem comes if all trading must be completed by a fixed end point in time, the resulting cost distributions can be highly left skewed if price trends are unfavorable. For example, if you are buying stock X and the price increases relative to the arrival price, most algorithms slow down the rate of executions to avoid the higher price. If the price goes up more, you slow down more. *If you run out of time with a significant portion of the trade not complete, you are now faced with either buying the remaining shares at the most unfavorable price of the entire process,* deferring the trade to a later date, or cancelling the remainder of the transaction altogether. Many methods to measure trading costs do not include the implicit cost of cancelled trades, and hence provide a downward biased estimate of the true costs.

Once an optimal trade schedule has been created we many need to further sub-divide the shares to be transacted in each period into a series of child orders typically of a few hundred shares each. In the ITOS implementation by ATM, an execution algorithm micromanages the sequence of small orders with a combination of real time liquidity models, randomization, active variation with market prices (as described in the preceding paragraph), and 'smart routing' to one or more potential trading venues which include the major exchanges, ECNs, and dark pools.

It should be noted that the analytical procedures described for both the pre-processor and the Northfield trade scheduler are entirely generic and do not specifically depend on the use of Northfield analytical models. In fact, ITOS even offers an alternative trade scheduling approach developed internally at ATM. However, we anticipate utilizing the most appropriate Northfield risk model for the trading needs of particular ITOS clients. These would normally be the US Short Term Model for US trades and the relevant "near- horizon" model for non-US transactions. Given that the time horizon of trading processes are rarely more than a couple days, adjustments will be made to the risk parameters to reflect the "fat-tailed" distribution of stock returns when measured over short periods. For more information on this aspect of the problem, please see our March 2008 newsletter, **http://www.northinfo.com/documents/285.pdf**. Similarly, the default source for market impact estimates will be Northfield market impact models, the free version of which has circulated to all risk clients since May of 2009 as described in our September 2008 newsletter, **http://www.northinfo.com/documents/311.pdf**.

(Portfolio-centric, continued on page 7)

(Portfolio-centric, continued from page 6)

Selected References

Schmidt, Thorsten. "Implied Risk Acceptance Parameters in Institutional Equity Trades", Proceedings of Northfield Seminar, http://www.northinfo.com/documents/246.pdf, June 2007

Wilcox, Jarrod. "Harry Markowitz and the Discretionary Wealth Hypothesis, *Journal of Portfolio Management*, 2003

Sharpe, William F. "Imputing Expected Security Returns From Portfolio Composition," *Journal of Financial and Quantitative Analysis*, 1974, v9(3), 462-472.

diBartolomeo, Dan. "Measuring Investment Skill Using the Effective Information Coefficient", *Journal of Performance Measurement*, Fall 2008.

Grinold, Richard C. and Mark Stuckelman. "The Value-Added/ Turnover Frontier," *Journal of Portfolio Management*, 1993,v19 (4), 8-17.

diBartolomeo, Dan. "Growth/Value/Momentum Returns as a Function of the Cross-Sectional Dispersion of Stock Returns", Northfield Research Conference, 2003, http://www.northinfo.com/documents/350.pdf.

Perold, Andre F. "The Implementation Shortfall: Paper Versus Reality," *Journal of Portfolio Management*, 1988, v14(3), 4-9.

Bertsimas, Dimitris and Andrew W. Lo. "Optimal Control Of Execution Costs," *Journal of Financial Markets*, 1998, v1 (1,Apr),1-50.

Northfield Website Update

Northfield President Dan diBartolomeo's latest blog entry where he discusses his insights on the Global financial crisis has been posted to the Northfield Website. Visit **http://www.northinfo.com/Essays/Blog23.pdf** to read the blog.

Dan Interviewed for Journal Article

Dan diBartolomeo was interviewed by David Spaulding in the Summer 2010 edition of the Journal of Portfolio Management. The interview has been posted to Northfield's website, **http://www.northinfo.com/emailimages/pm.pdf**.

Dan Interviewed for CFA "Take 15" Interview

Dan diBartolomeo was recently interviewed for a CFA "Take 15" interview where he discussed Applying Asset Liability Management to Private Wealth Management. The recorded interview can be accessed at: **http://www.cfainstitute.org/learning/products/multimedia/Pages/48638.aspx**

Northfield Speaking Engagements

Dan diBartolomeo spoke at the PRMIA conference in Montreal on December 1st. The topic was on Incorporating Liquidity Effects into Portfolio Risk Assessments.

Dan presented Equity Risk, Default Risk and Corporate Sustainability" at the New York QWAFAFEW on December 8th.

On February 24th, Dan will be speaking at the at the HI-FREQ TRADE conference in London. The topic will be on the "Parameterization of Trading Algorithms in Volatile Markets."

Nick Wade spoke at the Tokyo Quant Network Forum on November 18th. Northfield co-hosts this event with ClariFi. The topic was on Agent-Based Models.

Nick discussed Real Estate Risk Modeling at the FTSE/ EPRA/NAREIT seminar on December 8th.

Yasuhiko Nakase discussed Asset Allocation and Northfield's ART service at the Northfield Product Seminar in Tokyo on December 14th.

Northfield Now on LinkedIn

Northfield has now been established on the LinkedIn business oriented social networking site. The group currently has over 150 members. To become a member of the group, visit **http://www.linkedin.com/groupInvitation?groupID=2228261**

Northfield Partner Update

Northfield and Algorithmic Trading Management (ATM) announce the first release of the Integrated Trade Optimization Service (ITOS) targeted towards buy side trading desks and large pensions with significant amounts of internally managed assets. Please refer to Dan's piece on the first page of this newsletter for a discussion of our motivation for developing the product and the underlying approach taken and/or a product summary and implementation options at our web site at **http://www.northinfo.com/emailimages/atm.pdf**. Interested parties should contact Pete Hansen at ATM **phansen@atmanagement.com** for more details.

(Intra-Horizon, Continued from page 1)

volatility, both of which have increasing importance at shorter horizons.[3] In fact, the Basel Committee noted the importance of jumps and their omission from VaR in the Overview to the Amendment to the Capital Accord to Incorporate Market Risk (1996).

However, there is yet another dimension in which both tracking error and VaR are inadequate measures of risk, and that is in the time dimension. Both tracking error and VaR characterize the return distribution at the *end* of some investment period, and say nothing whatsoever about the return path followed by the investment *during* the period. Again, this omission was noted by the Basel Committee in the same 1996 document. Typically VaR is presented as a 10-day number, and tracking error as an annual number. However, neither measure tell us the probability or magnitude of the likely loss *within* that horizon, or "intra"-horizon.[4] We will use VaR-I to denote intra-horizon VaR.

The magnitude of intra-horizon risk is important in several situations:

1. Survival Risk: where there is a "floor" we cannot breach and remain in business. This is of particular concern to leveraged investors such as trading desks or hedge funds.
2. Monitored Investors: where the asset owner may remove the mandate due to poor performance within a particular horizon. If short-run performance is poor, the manager may be fired.
3. Valuation of Collateral: for example in securities lending or for capital adequacy regulations
4. Retirement/Endowment: for example, funding a liability intra-horizon

Luckily, this is not a new problem and has been the focus of great attention in the world of barrier-option pricing since active trading in barrier options began back in the 1960's. Of particular concern with barrier options is the probability of breaching a particular barrier during a holding period. Our concern is clearly related; given a probability level e.g. 1%, what barrier (VaR) might be reached?

To answer this question, we resort to a concept from statistics called "first passage probability," and also known as first hitting time. This is simply a group of measures that describe the probability that a particular level will be reached, or the expected time it will take to be reached, during a given horizon.

The first-passage probability for the Normal distribution is well known, including instances where there is drift in the mean.[5]

However, as previously noted, the Normal distribution is not always a suitable assumption. In cases where we need a more flexible distribution, we are forced to look for first passage solutions to more flexible distributions that allow jumps. Some examples would be Merton's "jump diffusion" model, or the two-sided pure jump model of Carr, Geman, Madan and Yor, or the finite-moment log stable model of Carr and Wu. These models can capture a much richer variety of distributions, but at the cost of some additional effort. There are few analytical solutions in these cases,[6] and typically we are forced to solve a partial integro-differential equation numerically by resorting to finite difference methods, binomial or trinomial trees, or Monte Carlo simulation.[7]

In the case of the Normal distribution without drift, we can derive the first passage probability and an expression for intra-horizon VaR quite easily by exploiting a relationship known as the principle of reflection. This simply states that for a given return path under the Normal distribution, in the absence of drift, there is an equivalent equally probable mirror-image return path.

For example, we can imagine a return path A that descends, breaches a barrier, and then recovers to end the period at some level "A" above the barrier. From the point at which A breaches the barrier we can draw a mirror-image path B that rises above the barrier mirroring A's fall below it, and then descends below it to land at some point "-A" below the barrier, mirroring A's rise above. Thinking about the probabilities for a moment, we can see that the path-dependent joint probability of breaching the barrier and ending the period at A is, by the principle of reflection, equal to the path-dependent joint probability of breaching the barrier and ending the period at –A below the barrier. In the absence of drift, things are symmetric about the barrier. The really fun bit is that the path-*dependent* joint probability of breaching the barrier and arriving at –A is equal to the path-*independent* probability of just ending at –A because to get to –A you have to breach the barrier *somewhere*. So by invoking the principle of reflection we have turned a path-*dependent* joint probability into a path-*independent* one.

We can then do a bit of probability maths:

Prob (breach barrier)

= 1 – prob(never breach)

= 1 – joint prob(end up above, and min always above)

= 1 – prob(end above) + joint prob(end above, and min <=barrier)

= prob(end below) + joint prob(end above, and min <=barrier)

(Intra-Horizon, Continued on page 9)

(Intra-Horizon, Continued from page 8)

We know from our discussion above that the second term in the last line; the joint probability of breaching the barrier and recovering to some level above the barrier is equal to the probability of ending at a symmetric point below.

Following the logic above, the probability of breaching a barrier b within a horizon T becomes:

$$\Pr(breach\_barrier) = N\left[\frac{\ln\left(B/S\right) - \mu T}{\sigma\sqrt{T}}\right] + \left(\frac{B}{S}\right)^{2\mu/\sigma^2} N\left[\frac{\ln\left(B/S\right) + \mu T}{\sigma\sqrt{T}}\right]$$

Note that I've used Girsanov's Theorem to convert a Normal distribution without drift into one with drift.

There are some very important implications that spring to mind immediately when examining that formula. First of all, the second term can never be zero or negative. That means intra-horizon risk is always greater than end-of-horizon risk. For example, even keeping the Normal distribution, prohibiting jumps, and setting the drift to zero: at 1% VaR-I = 1.107 * VaR[8]. At 2.5% it's 14.4% bigger, and at 5% it's 19.2% bigger.[9]

Note, as we will see later when we allow drift, non-Normal distributions and jumps, that multiplier can be considerably higher – in one instance up to 2.64[2]

Secondly, intra-horizon risk (probability) increases with time, whereas end-of-horizon risk (probability) decreases with time. This adds fuel to the argument against time-diversification. Various authors have argued that as the investment horizon increases (e.g. more years until retirement), the probability of loss declines and therefore an investor can afford to be more aggressive. In fact this is still taught in business schools, and forms a part of every investment suitability questionnaire that I have seen; "how old are you?". Samuelson (1963) argued against the benefits of time diversification, saying that although the probability of loss declines, the magnitude of the potential loss increases. When we consider risk from an intra-horizon perspective, we see that not only does the magnitude of the potential loss increase, but so too does the probability of loss.

Moving on, let's consider the impact of drift on VaR-I – but still with the Normal distribution. For comparative purposes, look at what happens when we introduce drift of 10% and 15% respectively. **(see table, top of next column)**

PROB	μ	VaR	VaR-I	VaR-I/VaR
5%	10%	$1.053\sigma\sqrt{T}$	$1.493\sigma\sqrt{T}$	1.417
2.5%	10%	$1.368\sigma\sqrt{T}$	$1.752\sigma\sqrt{T}$	1.281
1%	10%	$1.735\sigma\sqrt{T}$	$2.067\sigma\sqrt{T}$	1.191
5%	15%	$0.720\sigma\sqrt{T}$	$1.262\sigma\sqrt{T}$	1.753
2.5%	15%	$1.035\sigma\sqrt{T}$	$1.504\sigma\sqrt{T}$	1.453
1%	15%	$1.401\sigma\sqrt{T}$	$1.801\sigma\sqrt{T}$	1.285

We can clearly see that the introduction of drift has had a significant impact on both VaR and intra-horizon VaR, as suggested by the other work mentioned on "strategy risk." Instead of 10.7% bigger at 1% confidence without drift, VaR-I becomes 19.1% bigger at 10% drift, and 28.5% bigger at 15% drift. The multiplier for wider confidence bands (e.g. 2.5% or 5%) is worse still. For example, at the 5% confidence level and with 15% drift VaR-I is now 1.75 times the size of VaR – almost double.

Note that so far we have just stuck with the Normal distribution and have not allowed jumps, or stochastic volatility.

Various jump models have been proposed, but it is sufficient for our purposes just to review the empirical results of three of those:

- Merton (1976)
- Carr, Geman, Madan, Yor (2002)
- Carr and Wu (2003)

Jump models fall into two distinct groups; those that allow finite jumps (i.e. infrequent), and those that are called infinite jumps (i.e. lots)

As an aside, there is strong evidence for predictability of jumps, at least in the US markets. It seems that the VIX index can predict both jump arrival (timing) and size.[10] There is also evidence of jump clustering, and as a result we would tend to expect that the Merton (1976) model would not be as useful as the others since it cannot capture persistence in the data – it has no memory. However, at least empirically all three models discussed here seem to fit the data equally well – or at least not statistically significantly differently. The results are repeated from the Bakshi and Panayatov paper. **(see table below)**

(Intra-Horizon, Continued on page 10)

	Average VaR Multiples			Maximum VaR Multiples			Average VaR-I Multiples			Maximum VaR-I Multiples		
	JD	CGMY	FMLS	JD	CGMY	FMLS	JD	CGMY	FMLS	JD	CGMY	FMLS
S&P 500	1.21	1.33	1.44	1.33	1.37	1.69	1.60	1.39	1.94	1.77	1.50	2.08
FTSE	1.20	1.21	1.35	1.38	1.44	1.50	1.50	1.38	1.85	1.70	1.55	2.07
Nikkei	1.14	1.11	1.39	1.19	1.25	1.75	1.28	1.27	1.80	1.40	1.34	2.12
ATM Call	1.37	1.29	1.68	1.39	1.31	1.85	1.61	1.45	2.55	1.67	1.47	2.64

(Intra-Horizon, Continued from page 9)

These multiples are compared to standard Normal VaR:

- JD = Merton's jump-diffusion model
- CGMY is the two-sided pure-jump Levy model of Carr, Geman, Madan, and Yor
- FMLS is the finite-moment log-stable model of Carr and Wu

We can clearly see that

1. VaR with jumps is bigger than standard Normal VaR, although no doubt we suspected as much beforehand…
2. The choice of model makes a difference, although they all fit the data equally well. The FMLS model of Carr and Wu allows the fattest tailed distributions, and hence the most dramatic multipliers.
3. VaR-I is consistently greater than VaR, which we suspected based on the known result with the Normal, no drift case.
4. VaR-I can be more than *double* standard VaR, and this goes some way toward justifying the Basel multipliers.

At this point we have relaxed the Normal distribution assumption, allowed drift, and allowed for jumps in the return process under various different distributions. However, of course there are more things that we could try. Here are a few pointers for further reading:

1. Stochastic volatility in a VaR context - Eberlein, Kallsen, and Kristen (2002)
2. Two-dimensional PIDEs for models with both jumps and stochastic volatility - Feng and Linetsky (2006)
3. Leaning heavily once again on the barrier-option analogy we could also consider:
   - A risk *range* rather than a single loss level (double barrier)
   - A *time-varying* barrier or range
   - A barrier that varies as a function of volatility
   - wwE.g. Lo & Hui (2007), pricing double barriers where volatility, dividend yield, and the barriers are stochastic.

To summarize, intra-horizon risk is an important and neglected topic. If you are a buy-side portfolio manager you can probably implement and use Normal distribution formulation for first passage probability with drift given earlier in this note, since the central limit theorem has a surprisingly rapid effect, and even fairly concentrated portfolios rapidly converge toward Normal distributions over all but the highest frequency.[11] If you are a levered investor or a trading desk, then you will most likely need to resort to more flexible processes that allow fat tails and jumps in returns.

End Notes

[1] For example: tail-conditional expectation, worst conditional expectation. See Artzner, Delbaen, Eber and Heath (1997, 1999). Expected shortfall, CVaR see Acerbi, Tasche (2001)

[2] http://www.northinfo.com/Essays/strategy_risk_2010.pdf

[3] For evidence for jumps, see Bakshi, Cao, Chen (1997), Bates (2000), Anderson, Benzoni, Lund (2002). There is also evidence for large jump risk premia; Pan (2002)…but difficulty explaining market crashes with return jumps and diffusive volatility… For evidence for jumps in return and in volatility see Eraker, Johannes, Polson (2003). For two examples of models with return and volatility jumps see Duffie, Pan, Singleton (2000)

[4] This is a largely overlooked topic. There are a few good papers: Stultz (1996), Kritzman and Rich (2002), Boudoukh, Richardson, Stanton, Whitelaw (2004). However, these papers make restrictive assumptions: Brownian motion, no jumps, stationary volatility etc. One paper addresses the distributional issues and includes results for processes with jumps: Bakshi, Panayatov (forthcoming JFE)

[5] See for example Karatzas and Shreve (1991)

[6] Kyprianou (2006)

7 If you do resort to MC approaches, please take the time to read the excellent paper by Atiya & Metwally (2002) – they present a very fast approach, around 100 times faster than pure MC by leveraging a "Brownian bridge" to reduce the number of points that need to be calculated

[8] Feller (1971), 1% level.

[9] Boudoukh, Richardson, Stanton, Whitelaw (2004)

[10] Johannes, Kumar, Polson (1999)

[11] See e.g. Hlawitschka and Stern (1995)

References

- Acerbi, C., Tasche, D., 2001. Expected Shortfall: a natural coherent alternative to Value at Risk. Working paper.

- Anderson, T., Benzoni, L., Lund, J., 2002. Towards an empirical foundation for continuous time equity return models. Journal of Finance 57, 1239-84.

- Artzner, P., Delbaen, F., Eber, J., Heath, D., 1997. Thinking Coherently. Risk 10, 68-71

- Artzner, P., Delbaen, F., Eber, J., Heath, D., 1999. Coherent Measure of Risk. Mathematical Finance 9, 203-228.

- Atiya, A. F., Metwally, S. A. K., 2002. Efficient Estimation of First Passage Time Density Function for Jump Diffusion Processes. SIAM Journal on Scientific Computing 2002 (online version revised 2004)

- Bakshi, G., Panayatov, G., (forthcoming). First-Passage probability, Jump Models, and Intra-Horizon Risk. Journal of Financial Engineering (forthcoming)

- Basel Committee, 1996. Overview of the amendment to the capital accord to incorporate market risks. Working paper. Basel Committee on Banking Supervision.

- Bates, D., 2000. Post '87 Crash fears in S&P 500 futures options. Journal of Econometrics 94, 181-238.

- Boudoukh, J., Richardson, M., Stanton, R., Whitelaw, R., 2004. Max var: long-horizon value-at-risk in a mark-to-market environment. Journal of Investment Management 2, 1-6.

- Carr, Geman, Madan, Yor, 2002.

- Carr, Wu, 2003.

- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump diffusions. Econometric 68, 1343-1376.

- Eraker, B., Johannes, M., Polson, N., 2003. The Impact of Jumps in Volatility and Return. The Journal of Finance 58, 1269-1300.

- Feller, W., 1971. An Introduction to Probability Theory and Its Applications, Vol. II. John Wiley & Sons, New York.

- Johannes, M., Kumar, R., Polson, N., 1999. State Dependent Jump Models: How do US Equity Indices Jump? Online.

- Karatzas, I., Shreve, S., 1991. Brownian Motion and Stochastic Calculus, second edition. Springer (New York), Graduate Texts in Mathematics #113, 1991, 196-197.

- Kritzman, M., Rich, D., 2002. The mismeasurement of risk. Financial Analysts Journal 58, 91-99.

- Lo, C.F., Hui, C.H., 2007. Valuing double barrier options with time dependent parameters by Fourier series expansion. Published online February 2007.

- Merton, Reiner, Rubenstein, 1973

- Merton, 1976

- Jorion, P. 1997. Value at Risk. Chicago, IL: Irwin.

- Rosenberg B. and Guy J. "The Prediction of Systematic Risk" Berkeley Research Program in Finance, Working Paper 33, February 1975.

- Samuelson, P.A. 1963. "Risk and Uncertainty: A Fallacy of Large Numbers." Scientia, vol. 57, no. 6 (April/May): 108-113

- Stultz, R., 1996. Rethinking risk management. Journal of Applied Corporate Finance 9, 8-24

Tech Support Tip: Bayes-Stein Return Covariance (Return Shrinkage)

By Mike Knezevich

Frequent criticism of mean-variance optimization is the possibility of error maximization as an optimal portfolio is only as good as the inputs the user supplies (see DiBartolomeo). Northfield has introduced a series of new functionalities to decrease the impact of misspecified input data. This article discusses the Bayes-Stein adjustment in the Northfield Optimizer which uses a Bayesian approach similar to the approach included within the ART product (http://www.northinfo.com/documents/11.pdf). Function settings are located under the Estimation Error Adjustments node, where the adjustment is enabled and a Return Shrinkage weight specified.

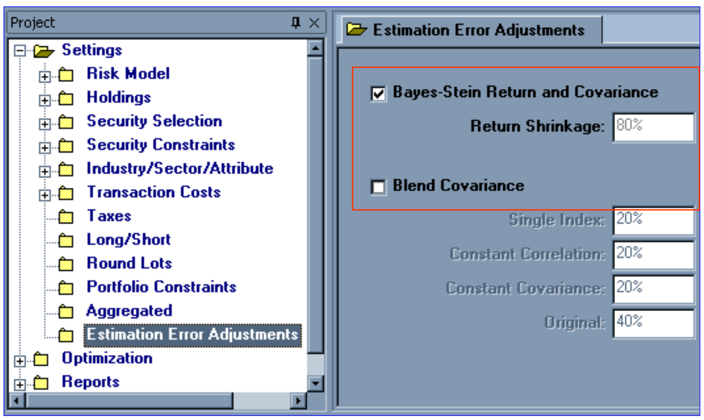

User supplied alphas are adjusted according to the expected return of the minimum variance portfolio created from the universe of assets. Additionally the risk model is slightly adjusted to account for changes in the underlying assumptions.

The Bayes-Stein adjustment is illustrated via an example which, for consistency, is similar to previous newsletter articles and may be found in the C:\NorthInfo\nisopt2008\samples2008\Estimation Error\Bayes-Stein directory. The buy universe consists of all the constituents of an equal weighted benchmark with corresponding alphas:

ID	Name	BnchWt(%)	Alpha
BBBY	Bed Bath & Beyond	25	-0.2032
LOW	Lowe's Companies	25	0
WEC	Wisconsin Energy	25	2.9751
XEL	Xcel Energy	25	7.4028

Table 1

Return Adjustment:

Returns are adjusted by squeezing the user supplied alphas to a prior return calculated from the minimum variance portfolio.

1) Begin by determining the minimum variance portfolio using the Optimizer then calculating the prior return. Recall that the minimum variance portfolio is the portfolio on the efficient frontier with the least amount of risk.

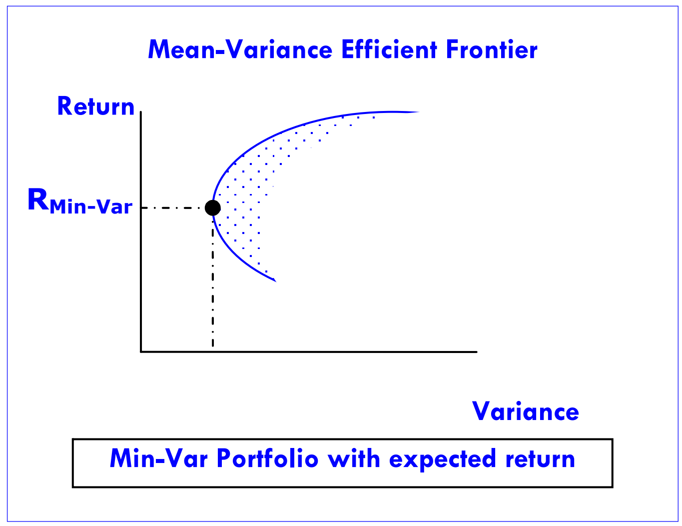

a) Simply run an unconstrained optimization with alphas against the benchmark to determine the minimum variance portfolio.

b) Weight the user supplied alpha according these optimal weights.

ID	Name	Opt Wt(%)	Alpha	Wtd Alpha
*$$$	CASH	0	0	0.0000
BBBY	Bed Bath & Beyond	11.099	-0.2032	-0.0226
LOW	Lowe's Companies	7.907	0	0.0000
WEC	Wisconsin Energy	46.308	2.9751	1.3777
XEL	Xcel Energy	34.686	7.4028	2.5677
				3.9229

Table 2

(Tech Support Tip, continued on page 12)

(Tech Support Tip, Continued from page 11)

The weighted average alpha of the minimum variance portfolio is the prior to which the user alphas will be squeezed and is the minimum return the user should expect given the inputs. In this example as calculated above, (R_{Prior})=**3.9229%**.

2) Adjust the supplied alphas ($R_{Expected}$) to the prior return (R_{Prior}) according to the user's Return Shrinkage setting.

   a. The Northfield default Return Shrinkage for the prior is used (w_{Prior}=80%).

   b. Supplied alphas are weighted as (1-w_{Prior})= 20%

   c. Calculate the Bayes-Stein Adjusted return:

$$R_{Bayes-Stein} = ((1-w_{Prior}) \times R_{Expected}) + (w_{Prior} \times R_{Prior})$$

For the example assets, the corresponding Bayes-Stein adjusted alphas are:

ID	$R_{Expected}$	20% of $R_{Expected}$	80% of R_{Prior}	$R_{Bayes-Stein}$
*$$$	0.0000	0.0000	3.1383	3.1383
BBBY	-0.2032	-0.0406	3.1383	3.0977
LOW	0.0000	0.0000	3.1383	3.1383
WEC	2.9751	0.5950	3.1383	3.7333
XEL	7.4028	1.4806	3.1383	4.6189

Table 3

Adjusted returns are more closely aligned to the prior due to the high Return Shrinkage rate. Assets which did not have supplied returns now have the prior (minimum variance portfolio) return.

Risk Adjustment:

Risk models are calculated using a sample mean; the return adjustment has effectively changed the mean increasing variability to the estimates. Factor and residual variances are adjusted to account for the increased variability.

1) The factor variance-covariance matrix is adjusted appropriately to reflect the adjusted mean.

$$\Sigma_{adjusted} = \Sigma + \delta_1 \Sigma + \delta_2 (1 \cdot 1^T)$$

Where:

Σ = Original Northfield supplied variance-covariance matrix.

δ_1 = Adjust the variance-covariance matrix based on the variation of the adjusted mean around the sample mean.

δ_2 = Adjust for uncertainty in model represented by a new factor (Est Err - Bayes-Stein Adj) in the matrix. All securities excluding the risk free asset has an exposure of 1.

The calculations of δ_1 and δ_2 are beyond the scope of this article, but are discussed in the referenced article by Jorion. Parameters for the example are calculated as δ_1 = .005 and δ_2 = 7.4501

Compare the original model variance (Σ) versus the adjusted model variance ($\Sigma_{adjusted}$) in the RiskModelTable:

Factor	Σ	$1+\delta_1$	$\Sigma_{adjusted}$
Beta	500.8080	1.0050	503.3350
Earnings/Price	5.1398	1.0050	5.1658
Book/Price	20.6965	1.0050	20.8009
Dividend Yield	6.9741	1.0050	7.0093
Trading Activity	8.4570	1.0050	8.4997
Relative Strength	56.3272	1.0050	56.6114
Log of Market Cap	7.9395	1.0050	7.9795
Earnings Variability	4.0625	1.0050	4.0830
EPS Growth Rate	7.0999	1.0050	7.1358
Revenue/Price	9.0285	1.0050	9.0741
Debt/Equity	7.5065	1.0050	7.5444
Price Volatility	15.8397	1.0050	15.9196
Retail Hard Goods	243.0780	1.0050	244.3045
Electric Utilities	123.2870	1.0050	123.9091
Est Err - Bayes-Stein Adj (δ_2)			7.4501

Table 4

2) Adjustments are also applied to residual risk at the asset level.

   a. The original residual variance (Resid Original) is multiplied by $\sqrt{(1+\delta_1)}$ to account for the adjusted mean.

   b. Assets have a unit exposure (EE_BSA) to the new Bayes-Stein Adjustment factor while cash is not exposed.

(Tech Support Tip,, Continued on page 13)

(Tech Support Tip, Continued from page 12)

These changes are reflected in the Maintable:

ID	Resid Original	sqrt(1 + δ_1)	Resid Adjusted	EE_BSA
*$$$	0.0000	1.0025	0.0000	0
BBBY	42.8509	1.0025	42.9589	1
LOW	47.4928	1.0025	47.6125	1
WEC	18.4637	1.0025	18.5102	1
XEL	20.3343	1.0025	20.3855	1

Table 5

Comparison:

Bayes-Stein adjustment's impact is illustrated by comparing optimizations of an initial cash portfolio against the equal weighted benchmark using the supplied alphas under different settings:

1) An unadjustment optimization where the Bayes-Stein is not enabled

2) A Bayes-Stein adjusted optimization with Return Shrinkage set to 80%.

Comparing the asset composition of the different optimizations note the unadjusted optimal portfolio takes greater bets in assets with higher returns while the Bayes-Stein adjusted optimal portfolio more closely aligns to the benchmark. This is not a surprise considering the alpha range has changed to (3.0977, 4.6189) from (-0.2032, 7.4028). XEL is still has the highest return and subsequent holding, but the impact has been muted

ID	Name	Bn Wt(%)	UnAdjusted	Bayes-Stein Adjusted
*$$$	CASH	0	0	0
BBBY	Bed Bath & Beyond	25	18.1506	23.6371
LOW	Lowe's Companies	25	19.6880	23.9429
WEC	Wisconsin Energy	25	2.9117	20.6047
XEL	Xcel Energy	25	59.2497	31.8153

Table 6

Comparing the utility decomposition of the different optimizations note the adjusted optimal portfolio has both lower expected return and risk:

	UnAdjusted		Bayes-Stein Adjusted	
	Return	Risk(v)	Return	Risk(v)
Factor	0	14.5	0	0.58
Stock Specific	4.44	80.11	3.72	3.19
Total	4.44	94.61	3.72	3.77
Tracking Error	9.73		1.94	
Portfolio Utility	3.49		3.68	

Table 7

Without complete confidence in the alpha and risk forecast, blunting the impact of estimation error can help avoid large blow ups in a manager's portfolio. Shrinking the return estimates to the minimum variance portfolio and introducing additional risk due to the change in the mean can effectively decrease the impact of outliers or misspecifided data. Bayes-Stein adjustment is one of the many new product features Northfield has recently added to address the issue of error maximizing in mean variance optimization.

For further inquiries, contact Technical Support in Boston: support@northinfo.com or call 617.208.2080. European clients can contact: support-europe@northinfo.com or call +44-(0)-20-7801-6260. In Asia, call +81(0)3 5403 4655 or +61(0)2 9238 4284 or support-asia@northinfo.com.

References:

DiBartolomeo, Dan. "Portfolio Optimization: The Robust Solution", Prudential Securities Quantitative Conference, December 1993. http://www.northinfo.com/documents/45.pdf

Jorion, Phillippe. "Bayes-Stein Estimation for Portfolio Analysis", *Journal of Financial and Quantitative Analysis*", 1986, September V21(3), 279-292.

Warrick, Sandy. "Estimating Returns for Asset Allocation", Northfield Asset Allocation Seminar, September 2004. http://www.northinfo.com/documents/26.pdf

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
Shakespeare House
168 Lavender Hill
London, SW11 5TG
Phone: +44-(0)-20-7801-6260
Fax: +44-(0)-20-7801-6261

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 INFORMATION SERVICES, INC.

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