July 2003

Northfield News

Quarterly Newsletter for the Friends and Clients of Northfield Information Services

Open Performance Attribution Released to Users

By Dan diBartolomeo, Robert Kelley and Tracy Licklider

On June 30th, we released the first production version of our totally new Open Performance Attribution system. The release was on schedule, and was subsequent to an extensive period of beta testing with a number of clients whose efforts and cooperation are much appreciated.

This "Open" application replaces our long-standing Performance Attribution system with greater flexibility, speed and more detailed analysis than exists in any other commercially available system. The most crucial feature of the new system is the ability to operate with any multiple factor model of security behavior. All Northfield models can now be used for attribution. In addition, models from other vendors and user-defined models can be fully supported. This will allow investment firms to custom tailor their attribution process to specific investment products and themes. Users can define a variety of factor types (e.g. continuous variables such as beta, or P/E, binary variables such as sector membership).

It is also possible to combine an existing Northfield model with a user defined variable such as an alpha forecast. This will allow users to have a convenient way of getting detailed analysis of the effectiveness of their alpha estimates, net of the impact of the prescribed set of risk factors. Important distinctions between "bottom-up" and "top-down" strategies can be examined in this fashion.

(Continued on page 2)

Active Risk Budgeting Using Northfield Systems

By Dan diBartolomeo

One of the practices gaining popularity among plan sponsors and other multiple-manager funds is the concept of "risk budgeting." In traditional mean-variance asset allocation methods, investor's make assumptions about the expected returns, variances and correlations among the assets in which they intend to invest. However, if these assets are themselves managed funds as opposed to passive indices representing asset classes there is an additional layer of complication. While we may be confident in our forecasts as to the expected behavior of asset classes, we may be far less confident in our forecasts that particular actively managed funds will outperform their respective benchmarks.

Risk budgeting is an approach to address this issue. There are many potential variations on the theme of risk budgeting but the key feature is to separate the decisions about asset allocation from the decisions about how aggressive we choose to be in each actively managed asset class. The investor first decides upon an allocation of their capital to passive

(Continued on page 5)

Inside this Issue

Special points of Interest:

- London Office Client Services added
- New Enhancements to the Northfield Everything Everywhere Risk Model
- Risk Budgeting Using the Northfield System

Newport Summer Seminar Wrap Up

Tennis Hall of Fame • Newport, RI • June 6, 2003

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 a sold-out audience of Northfield clients and friends.

International Tennis Hall of Fame

The agenda consisted of 6 presentations including: Solutions for Improving Investment Process efficiency, Bayesian Analysis of Stochastic Beta, Multiple Factor Performance Attribution, Reconciling Fiduciary and Social Responsibility, Fund Manager Behavior and the Tournament Mentality, and Real World Portfolio Optimization for Taxable Accounts.

As is customary, the seminar coincided with the USA Professional Championship of Court Tennis. Following the seminar, attendees viewed a court tennis demonstration, and then a Semi-Final Match between Tim Chisholm of the US and Nick Wood of the UK (Chisholm won in three straight sets). Court Tennis, or "real tennis" is the medieval sport that is the progenitor of all modern racquet sports.

After tennis on Friday evening, everyone enjoyed a relaxing dinner party held at Johnnie's Atlantic Beach Club and Pavilion in nearby Middletown. A three piece band and vocalist provided the evening's entertainment. Complete proceedings have been posted to our website at http://www.northinfo.com/papers.

(Open Performance, Continued from page 1)

In keeping with our existing approach to attribution, the new system provides both stratification style univariate attribution (Brinson-Fachler) and multivariate factor-based analysis. Having both popular forms of attribution interwoven in the same reports allows users to gain understanding of investment effects that would be hidden if either methodology were used alone. Consistent with the "generic" nature of the system, existing reports have been modified slightly to no longer require that models fit into the framework of the Capital Asset Pricing Model. In the new reporting scheme, users can also aggregate factors into groups of their own choosing so as to distinguish between the effects of value and growth factors, or equity and fixed income aspects of a balanced portfolio. All reports can be viewed and printed through an entirely new Windows interface, or easily exported to Excel.

The risk reporting aspects of the system have been enhanced. All analytical procedures regarding risk estimation are now entirely consistent with the methods and conventions used in the Open Optimizer. We have also expanded the analysis of the role of cash in portfolios, including cash in some reports and excluding it from others as seems most analytically appropriate for that report content.

An important enhancement in the "ease of use" aspect of the system is that many of the underlying data files are now shared with our Optimizer product. All files are in ASCII format and can be easily edited from within the new system's Windows interface, or by using an external spreadsheet program. This should dramatically reduce the user overhead in maintaining data files with respect to user edits such as additions to the data set for privately placed securities or to account for an identifier change between data updates. Another improvement in user friendliness relates to exception handling, where the reporting of any data irregularities or conflicts with the problem set up has been extensively improved. A capability has also been introduced to conveniently allow for "what if" analysis of portfolios that are held constant through time.

The Open Performance System also supports flexible definitions of analytical periods. Users can define models different observation periods such as daily or monthly. Northfield models are delivered on a monthly basis and will continue to be so. We believe that much conventional statistical analysis used for performance attribution is improper when daily data is involved. This issue is the subject of a research paper by Dan diBartolomeo, "Frequency of Observations in Performance Attribution: Just Because We Can Doesn't Mean We Should", that is forthcoming in the *Journal of Performance Measurement*. This paper also

(Continued on page 3)

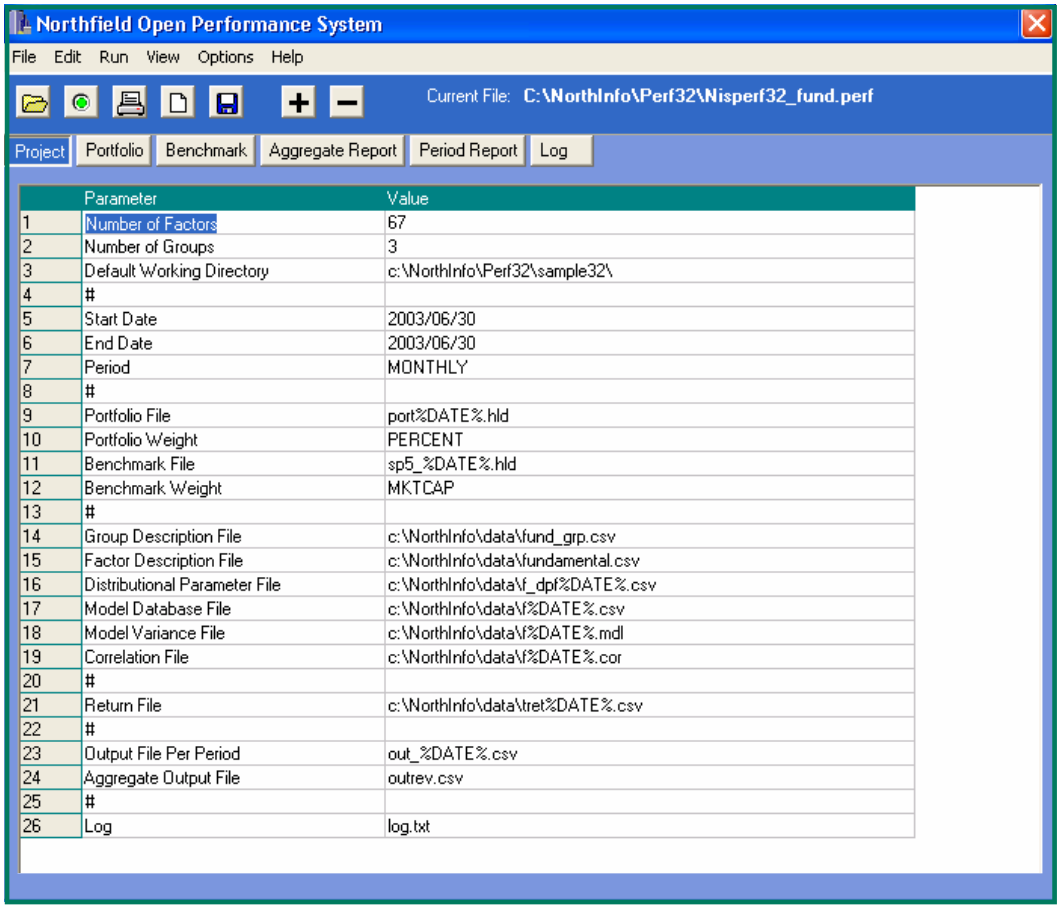

	Parameter	Value
1	Number of Factors	67
2	Number of Groups	3
3	Default Working Directory	c:\NorthInfo\Perf32\sample32\
4	#	
5	Start Date	2003/06/30
6	End Date	2003/06/30
7	Period	MONTHLY
8	#	
9	Portfolio File	port%DATE%.hld
10	Portfolio Weight	PERCENT
11	Benchmark File	sp5_%DATE%.hld
12	Benchmark Weight	MKTCAP
13	#	
14	Group Description File	c:\NorthInfo\data\fund_grp.csv
15	Factor Description File	c:\NorthInfo\data\fundamental.csv
16	Distributional Parameter File	c:\NorthInfo\data\f_dpf%DATE%.csv
17	Model Database File	c:\NorthInfo\data\f%DATE%.csv
18	Model Variance File	c:\NorthInfo\data\f%DATE%.mdl
19	Correlation File	c:\NorthInfo\data\f%DATE%.cor
20	#	
21	Return File	c:\NorthInfo\data\tret%DATE%.csv
22	#	
23	Output File Per Period	out_%DATE%.csv
24	Aggregate Output File	outrev.csv
25	#	
26	Log	log.txt

Open Performance Attribution Project Screen

(Open Performance, Continued from page 2)

proposes a method to reconcile the results of using "monthly buy-and-hold" observations with daily observations. We anticipate incorporating this new methodology into a future release of the system.

The ability of the system to handle production runs has been dramatically improved. An activity log now records all problems processed and records any irregularities during a particular analysis run, while allowing the remainder of the batch process to proceed. The standard product installation includes a "command line" version that can be run under DOS or Windows. The underlying performance attribution engine can also be called from programs written in Visual Basic, Java, C++, C#, and most other modern programming languages. The Open system is also an order of magnitude faster than the old system making short work of even large production tasks.

The next major release of the Open system will include a set of tools to assist users in defining their own factor models. Assuming a user provides periodic files of factor exposures for their universe of securities, the system will be able to do cross-sectional regression analyses to compute factor returns for each time period. The system will also compute its own time series factor variances and factor correlations so that user-defined models can be used in the system's risk reporting as well.

Upcoming releases will also incorporate a variety of new reports and analytical methods. A new "marginal contribution" has been designed to show whether securities positions that are contributing the most risk to a portfolio are also contributing the most return. From this information we can calculate a "realized information coefficient" to jointly estimate whether a manager is both good at forecasting security returns and weighting their positions so as to take best advantage of their forecasts. Other upcoming reporting items include a non-parametric measure of diversification (the J statistic) and a new analysis of portfolio turnover. Finally, we will be reporting a further delineation of factor return impacts to distinguish between the effects of consistent factor exposure bets (tilts) and the effects of factor exposure timing.

From a software perspective, upcoming releases will include a rich graphics library for preparing a variety of output charts, compatibility with UNIX and LINUX, and a DLL version that will allow users to incorporate our calculation engine directly into in-house systems.

Major Enhancements Made to the Northfield Everything Everywhere Model

The July 2003 monthly update marked a major improvement in Northfield's Everything Everywhere risk model. This change is a result of our ongoing commitment to our clients' needs and maintaining the highest level of accuracy in our analytical tools.

Clients have reported concerns that the original methodology failed to address the difference between major fixed income markets, and also failed to address the sometimes quite substantial volatility differences between major fixed income markets. To address these concerns we have changed the pricing methodology we use, and also added an adjustment to the sensitivities we calculate.

The improvement concerns the framework used to price fixed income securities, and to determine their sensitivity to changes in the term structure shape. In the original version of the model, each fixed income security was priced using the dominant yield curve for the region in which that bond was issued. For example, a French bond would be priced using the dominant yield curve for the "Continental Europe" region; a UK bond using the dominant curve for "English Speaking Countries" (the US curve) and so on. Whilst this regional breakdown works particularly well for equity securities, it does not apply so well to fixed income securities. In our new and improved model, fixed income securities from the five largest issuing markets will be priced based on their own individual term structures - US dollar, British pound, Japanese yen, Swiss franc, and Euro. Historical yield structures are constructed using yield data from the Bank of England, Bank of Japan, Bank of Switzerland, and the European Central Bank. All other instruments are valued using a treasury yield blended from these five curves, and adjusted using exchange rate volatility.

To go into the process in more detail: Yield curves for securities issued in currencies other than the five major ones are constructed utilizing currency exchange rate volatility. Specifically, this is done by weighting the yields from the five major currency curves by the inverse of exchange rate volatility. For instance, for constructing the yield curve based in Thai baht (THB) we will use as weights for the respective yields the inverse of the volatility of the exchange rates of respectively the USD/THB, GBP/THB, JPY/THB, CHF/THB,EUR/THB. This method is based on the premise that if term structures in two different currencies are similar to each other the exchange rate volatility on those will be lower. If there were minimal volatility between the two currencies (e.g. USD and THB) this would mean that the Thai bath yield curve will be identically shaped as the US curve and the two will differ just by a spread (this spread is captured as an element of risk as has been presently done in the calculation of Option Adjusted Spread for each security). Once we have constructed yield curves based in all denomination currencies covered by the model, we price and calculate risk parameters for each fixed income security off its particular curve.

The term structure factors reported by the model and the security exposures to those factors have also undergone a related modification. The figures shown in the latest .MDL file will be the factor values for the Global Yield Curve. The Global Yield Curve is a construct based on the five major currency yield curves each weighted by the proportion of outstanding dollar amounts of treasury securities traded in the particular currency. The Global Yield Curve factors serve as a pivoting point for implicit calculation of risk parameters for both the individual security risk exposures as well as the factors of the yield curves based in all of the thirty-seven currencies. For this purpose, each security exposure is scaled by a ratio, which is given as the volatility of the respective factor from the respective currency's yield curve divided by the same factor from the Global Yield Curve. For example if a particular US bond has an exposure to term structure factor 2 (curvature) of 0.03 and the volatility of the same factor in the USD denominated yield curve is 1.25 more volatile than the same factor in the Global Yield Curve then the reported exposure in the EE data file will be 0.0375 (0.03 * 1.25). This implicitly makes sure that the variance the factor from each currency's yield curve is correctly utilized in the risk equation, while allowing the convenience and simplicity of handling the factors for just one term structure - the Global Yield Curve.

We believe that these two significant changes will go a long way toward resolving the concerns of our clients that our approach truly captures the risk of fixed income instruments as best it might.

If you have any questions or concerns regarding the changes, please contact our Technical Support department at 617.208.2080, support@northinfo.com.

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

(Active Risk Budgeting, continued from page1)
asset classes using whatever methodology that they choose. Once the asset class weights are determined, we then can separately take up the issue of how aggressive we would prefer our active managers to be within each asset class.

The basic risk budgeting procedure is to first select a limit on how much active risk we would like to take within our multiple manager fund overall. We then carry out an allocation procedure to determine how much active risk we want to take in each asset class. For example, we might want to have a passive total return volatility of 8% per year in our allocation among passive asset classes, and then have a tracking error (volatility of active return) of 2%.

Investors often make additional simplifying assumptions in order to make the mathematics of the problem more easily solved. Another key assumption is that all of our estimated parameters are known with certainty, so the problem can be solved via classic Markowitz optimization. This is a paradoxical choice because if we were sure that our forecasts of future return distributions were accurate, there would be no need to segment the problem into a passive asset allocation and an active risk budget. We could merely compute the expected return distribution of each actively managed fund (combining the market and active risks and returns within each asset class) and do a traditional asset allocation. Often the assumption of zero transaction costs is also made. In practice, restructuring active portfolios to have more or less risk, moving capital between active managers, or moving capital between asset classes all involve substantial costs.

Three more simplifying assumptions are frequently seen. The first simplifying assumption is that active risk and asset class risks are uncorrelated. However, this would clearly not be true in cases where the "active" strategy was merely to increase or decrease systematic risk within an asset class. A portfolio of high beta stocks will show positive correlation between the active return and the market return, while a portfolio of low beta stocks will show negative correlation. A second simplifying assumption is that there is no upper bound on the total volatility of the entire fund. The total volatility comes from three sources: the volatility of the passive asset classes, the tracking errors of the active managers of the active funds around the asset class benchmarks and the covariance between the active returns and the market returns (which we may choose to assume is zero as per the prior paragraph). If there is a limit on the total volatility, then the optimal passive allocation and the optimal risk budget must be solved jointly so that the total risk fits within the overall limit. A final assumption is that there are no upper bounds on the amount of tracking error that can be budgeting to a particular asset class. Obviously, there are practical limits on how aggressive an active manager can be.

Using Northfield's Open Optimizer or PACO Asset Allocation systems, active risk budgeting problems may be solved in a fashion that eliminates the need for any of these simplifying assumptions. The PACO system includes a Bayesian re-estimation process from Jorion ("Bayes Stein Estimation for Portfolio Analysis", *Journal of Financial and Quantitative Analysis*, 1986) that automatically adjusts the investor's return, volatility and correlation expectations for the impact of forecast uncertainty. The Open Optimizer includes the patented NFA resampling technique that uses a Monte-Carlo like simulation method to explicitly estimate and compensate for the uncertainty around portfolio properties such as tracking error. Both systems allow for inclusion of expected transaction costs and a variety of constraints on active risk for individual asset classes and combinations of active risk across asset classes. In addition, both systems allow for "zero capital" strategies such as currency overlays or tactical asset allocation (relative to the strategic asset allocation) to be included.

Let us now undertake an example. In keeping with the way most investors approach risk budgeting, we will presume that active returns are uncorrelated with market returns (no constant systematic bets) and that there is no limit on total fund volatility. In our fund, we have four asset classes with actual allocations:

Domestic stocks	40%
International stocks	10%
Bonds	10%
Cash	40%

In addition, we have a currency overlay manager that can speculate in the currency futures market with notional positions equal to 25% of the fund assets. Finally, we have a tactical asset allocation overlay manager who can adjust the asset allocation for the entire fund (100%) using index futures.

Assume we are mandated to have 2% tracking error level for the overall fund. Our job is to decide the ideal level of tracking error for each of the actively managed portfolios within the whole fund. This decision is conditional upon our beliefs as to the skills of each active manager and the extent to which we believe the active returns for each asset class will be correlated with the active returns for the other asset classes. We think our active managers are capable of the following information ratios (active return/tracking error), and that their active returns will be correlated in the following fashion: (See Table 1 on the following page)

(Continued on page 6)

Table 1

	IR	Correlations					
Domestic stocks	0.3	1.0					
International stocks	0.7	0.3	1.0				
Bonds	1.0	0.2	0.1	1.0			
Cash	0.0	0.2	0.0	0.1	1.0		
Currency	1.0	0.1	0.0	0.1	0.0	1.0	
TAA	0.5	0.0	0.0	0.0	0.0	0.0	1.0

Our active bond manager has indicated that they feel it impractical to have a tracking error greater than 3%. Our domestic stock active manager has indicated that 6% will be their upper limit on tracking error.

The key to our process is to realize that having an active manager with a tracking error of 4% is the same risk as having an active manager with a 1% tracking error, whose portfolio has been levered four times through the use of borrowed capital (shorting index futures on the benchmark so we only lever the active risk). We can now proceed from the assumption that each active manager has tracking error equal to 1.0. It is our job to figure out what is the ideal "leverage" to be associated with each active manager. This gearing ratio will be numerically equivalent to the ideal tracking error for each manager.

Since we assume a zero information ratio for the cash asset class, we wouldn't want to pay an active manager to carry out this investing. As such, the tracking error will be zero, as opposed to one for our other idealized managers. So this asset class drops out of the problem. Given tracking errors of one, the expected active return for each manager is simply equal to their information ratio. So we now have the following asset allocation problem.

Table 2

	Init	Min	Max	Ret	SD	Correlations					
Domestic stocks	40	0	240	0.3	1.0	1.0					
International stocks	10	0	NA	0.7	1.0	0.3	1.0				
Bonds	40	0	120	1.0	1.0	0.2	0.1	1.0			
Currency	25	0	NA	1.0	1.0	0.1	0.0	0.1	0.0	1.0	
TAA	100	0	NA	0.5	1.0	0.0	0.0	0.0	0.0	0.0	1.0
Slack	-115	NA	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

The initial weight column is set to the amount of portion of the fund that is subject to this particular form of active management. The minimum weight column illustrates that we could choose to go to a zero tracking error (passive management) in each asset class as appropriate. The maximum weight class reflects the upper bounds on tracking error as specified by our domestic stock manager (240/40 = 6) and our bond manager (120/40 = 3). All the standard deviations of the active management "assets" are set to one, and returns are equal to the expected information ratios. In that the initial weights add to more than 100%, we introduce a "slack" variable that is the mathematical representation of shorting the benchmark index to allow for leveraging of the active management. Since the slack variable represents taking a position in the benchmark it has no active return or risk by definition.

The asset allocation problem specified above can be routinely solved in either the Open Optimizer or the PACO Asset Allocation system for the desired fund level volatility of 2%. As previously noted, both systems can augment this basic computation with additional complexities such as trading costs, complex constraints, and explicit handling of estimation error in our forecasts. Once the problem has been run, the ideal tracking error for each actively managed asset class would be given by the "optimal weight" divided by the initial weight for that class.

By representing each asset class as two assets, the passive component and the active component, we can eliminate the need to assume that active returns are uncorrelated with passive market returns. Similarly, this augmented approach can also provide the "joint" solution of the risk budgeting problem and passive asset class allocation as required when bounds on total fund volatility exist.

London Office Client Support

The London office is expanding as a result of the significant growth in Northfield's European client base. We are pleased to announce that Christine Milne has joined Rupert Goodwin to provide client services, including technical support and training.

Christine recently graduated from the University of Cambridge with an upper second Natural Sciences BA in Materials Science and Metallurgy. Prior to joining Northfield, Christine broadened her horizons by working at Citigroup Asset Management and in an orphanage in Ghana.

Christine Milne - European Client Services, can be contacted by e-mail, christine@northinfo.com or at the London office on +44 (0)20 7801 6260.

Northfield Staff Speaking Engagements

Northfield President Dan diBartolomeo spoke at the Separately Managed Accounts Conference in Sydney Australia, on July 23rd. The topic was on the Management of Taxable Accounts. For more information, go to http://www.orkestrate.com.au/events/sma_2003/

On September 8th, Dan will be speaking at the Alpha Strategies Conference, in Oxford, England. The topic will be on Modeling Convertible Bonds Among The Everything Everywhere Model. Dan will also be speaking at the Green Mountain Summit on Investor Responsibility on September 16th. The topic is still unannounced.

On September 23rd, Dan will be presenting at the QWAFAFEW meeting in New York, the topic will be on Enhanced Index Funds. He will also be speaking at the Financial Management Association Annual Meeting in Denver, Colorado on October 9th. The topic will be Manager Compensation and Seasonal Variation in Fund Risk Levels. Visit http://www.fma.org/Denver for more information.

Northfield's Emilian Belev gave a presentation at the Toronto QWAFAFEW meeting on July 15th. The topic was on "Convertible Bonds Valuation and Risk"

Northfield Sponsors Sailing Team

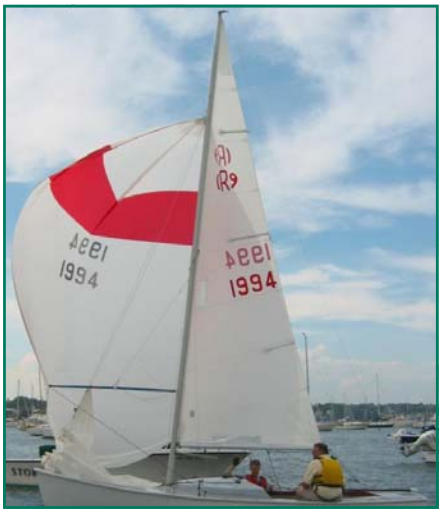

This summer, Northfield entered the 2003 Courageous Corporate Challenge. The Courageous Sailing Foundation, provides free sailing instruction for the entire summer for any child who lives in Boston or the nearby towns.

In recent years, a number of very well known sailors have participated in this event, including the 1984 Olympic medalist Steve Benjamin and Steve Ullian from Deutsche/Scudder, the J22 World Champion. Despite the fact that neither of them has sailed much for the last 10 or more years, Northfield's Sandy Warrick and Richard Pearce shook off of the rust and ended up tied for first on the first (practice) night. But we had heard a rumor that Nick Trotman of Adam, Harkness, Hill would be showing up the following week, he was a 1994 Collegiate All American and the 1998 505 World Champion.

Joining us on four nights was Steve Cucchiaro, who has been helping us develop Paco's new Tactical Asset Allocation module. Although Steve's tactical ability at asset allocation is how he makes his living, his tactical sailing ability is probably better known; he was a college 1973 All American (and Sandy's roommate) at MIT and has won national and Pan Am championships.

Northfield won three of the nights (honesty requires pointing out that those were the nights that Steve joined us), and Adam Harkness Hill won three nights and the cumulative scoring. But the real winners were the kids of Boston – the Challenge provided for sailing lessons for more than 200 children of Boston. Richard and Sandy are planning their fall, winter frostbite and spring practices to prepare for next summer.

Northfield INFORMATION SERVICES, INC.
www.northinfo.com

Boston Office
184 High Street, 5th Floor
Boston, MA 02110
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 5TF
Phone: +44-(0)-20-7801-6260
Fax: +44-(0)-20-7801-6261

Northfield News is a publication of Northfield Information Services, Inc., 184 High Street, 5th fl., Boston, MA 02110. 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