



June 9, 2008

Galen Burghardt

galen.burghardt@newedgegroup.com

Jennifer Cai

jennifer.cai@newedgegroup.com

Bin Chen

bin.chen@newedgegroup.com

Lianyan Liu

lianyan.liu@newedgegroup.com

Correlations and holding periods:

The research basis for the AlternativeEdge Short-Term Traders Index

Introduction (This we believe)

The work we have reported in *Superstars versus teamwork* and in *There are known unknowns* has persuaded us that the most reliable tool for improving a portfolio is diversification when correlations are low. It is, in short, the only reliably free lunch available in the world of investments. For one thing, correlations seem to be predictable while Sharpe ratios are not. For another, uncertainty about the true values of return distributions seriously hamper efforts to improve on a simple volatility weighted portfolio of low correlation assets.

At the same time, our ongoing conversations with managers and investors have produced a mountain of anecdotal evidence that short-term traders' returns should exhibit low correlations – not only with the returns of nearly any other investment alternative but with one another's returns.

For these reasons, we decided to tackle the problem of constructing an index of short-term traders' performance. We were confident that the effort would be rewarded by highlighting an investment space with very desirable correlative properties.

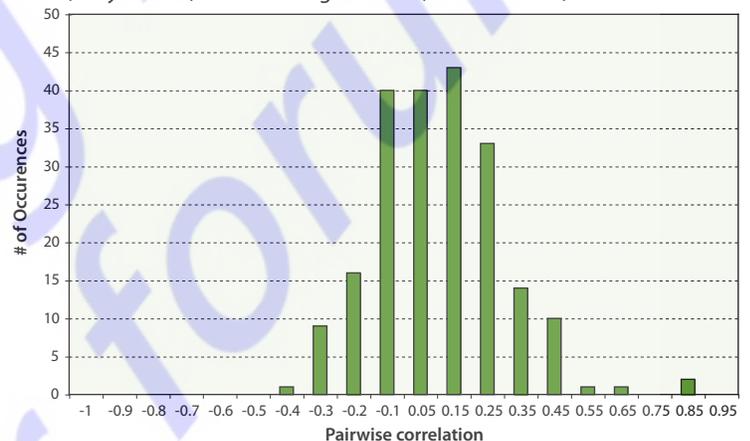
The purpose of this note is to discuss the properties of the index and various strategies for replicating the index. In the process, we briefly describe:

- key findings of *Superstars versus teamwork* and *There are known unknowns*
- construction of the index (STTI)
- correlations of STTI returns among themselves and with others'
- why the correlations are low
- why there are not more investors in this space
- some challenges in replicating the index

Review of previous research

Our confidence in correlation as an organizing principle for constructing an index stems mainly from the research we reported in *Superstars versus teamwork* (Newedge, May 4, 2007). Everyone knows, as shown in Exhibit 2, that low correlation is valuable because it allows for substantial risk reduction for any given level of expected returns. What is more important, though, is that correlation appeared to be predictable while past performance as measured by Sharpe ratios was not. Exhibit 3 provides an example of the relationship between historical and realized implied correlations within 10-manager

Exhibit 1
Distribution of pairwise correlations among 21 STTI components
(daily returns, 6/1/07 through 5/30/08, mean = 0.104)



AlternativeEdge series

The Alternative Edge series of research notes is a collaboration between Galen Burghardt and the Alternative Investments Group (AIG), which is led by Leslie Richman and Brian Walls. The AIG is committed to open and informed dialogue between investors and money managers. To that end, we endeavor to provide our clients with probing and objective research into questions that matter to us all. For a compilation of our research notes and information on the Alternative Edge CTA Indices, please go to newedgegroup.com.

The reader is advised that futures and options are speculative products and the risk of loss can be substantial. Futures spreads are not necessarily less risky than short or long futures positions. Consequently, only risk capital should be used to trade futures. The information contained herein is based on sources that we believe to be reliable, but we do not represent that it is accurate or complete. Nothing contained herein should be considered as an offer to sell or a solicitation of an offer to buy any financial instruments discussed herein. All references to prices and yields are subject to change without notice. Past results are no indication of future performance. Any opinions expressed herein are solely those of the author. As such, they may differ in material respects from those of, or expressed or published by or on behalf of, Newedge Group or its officers, directors, employees or affiliates. © Newedge Group, 2008

Exhibit 2
Risk reduction through diversification

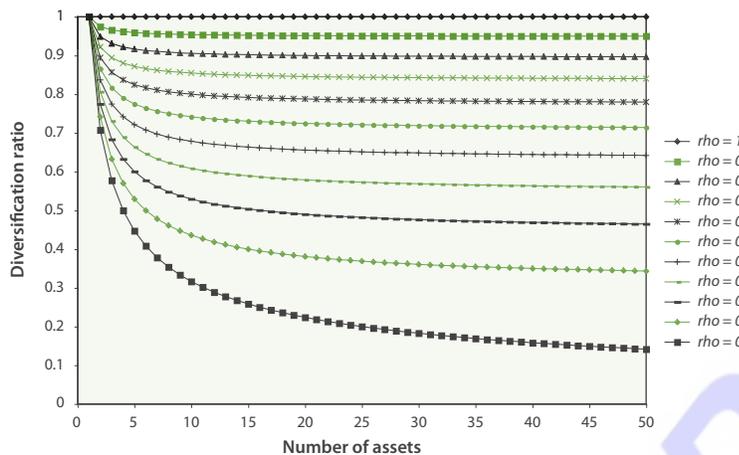


Exhibit 3
Historical and realized implied correlations for portfolios of 10 managers
(Historical 1995 - 1997, realized 1998)

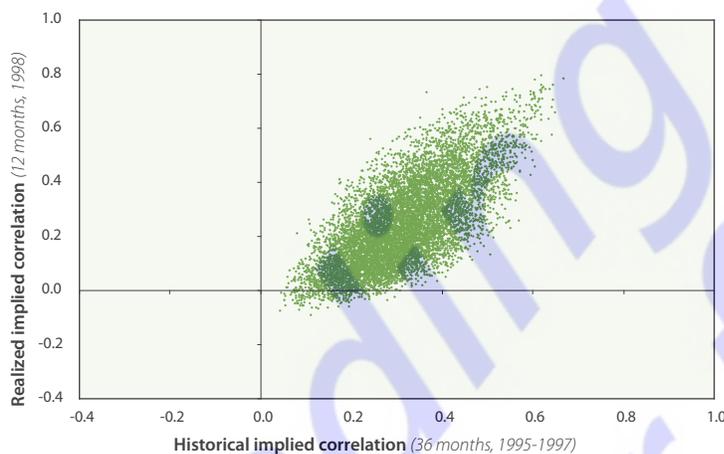
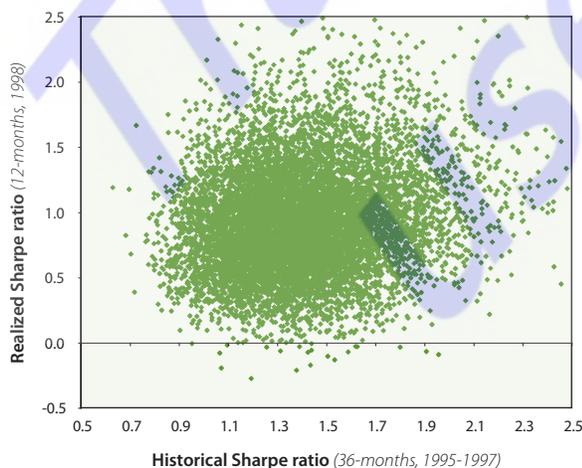


Exhibit 4
Historical and realized Sharpe ratios for portfolios of 10 managers
(Historical 1995-1997, realized 1998)



CTA portfolios for a look-back period of three years and a look-forward period of one year. In contrast, Exhibit 4 shows no relationship at all between historical and realized Sharpe ratios for the same portfolios. In these examples, the look-back period was 1995 through 1997, while the look-forward period was 1998. We conducted the same experiment for every similar set of years through 2005 and found roughly similar results.

The upshot of all this was that when we ran a horse race between portfolios of managers selected for their contributions to the portfolio and portfolios of managers selected for their individual performance, we found that the teamwork portfolios significantly outperformed the superstar portfolios. In fact, what we found, as shown in Exhibit 5, was that the superstar portfolios turned in what would have been average performance for the set of managers from which we were drawing portfolios, while the teamwork portfolios produced overall Sharpe ratios that were well up in the distribution of what one could have achieved by constructing portfolios at random.

The peculiar shape of the scatter in Exhibit 3 also produced a surprising insight into the difference between estimates of implied or average correlations within a portfolio and what one expects to find when estimating pairwise correlations. In particular, when correlations are truly low, the distribution of correlation

Exhibit 5
Distribution of Sharpe ratios for randomly selected portfolios

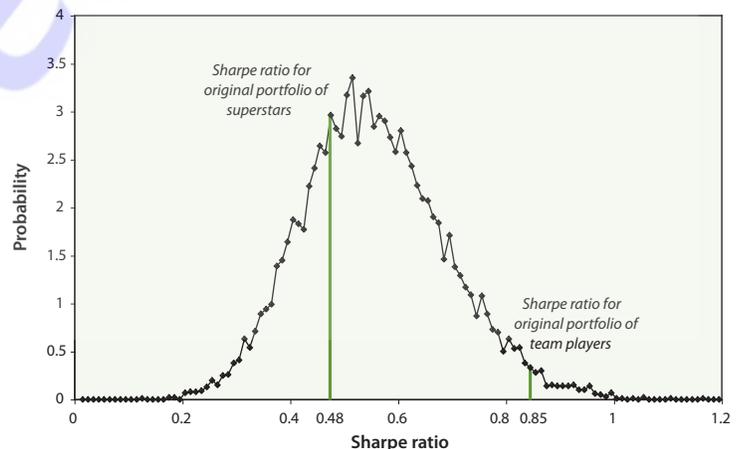


Exhibit 6
Improvement in estimates of implied correlation as the number of assets increases

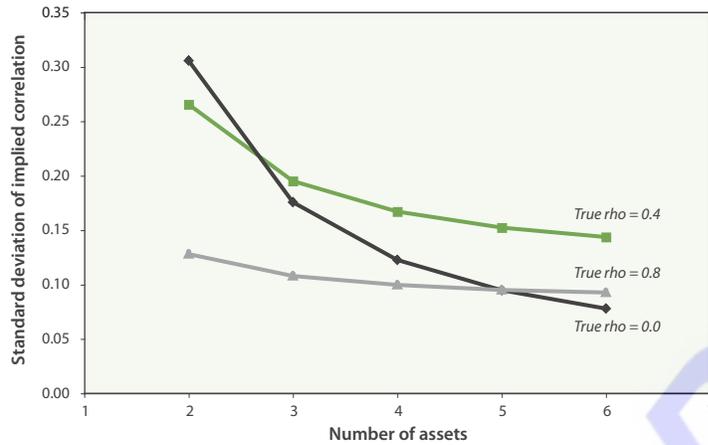


Exhibit 7
1-standard-deviation bands for the Sharpe ratio (12 monthly returns)

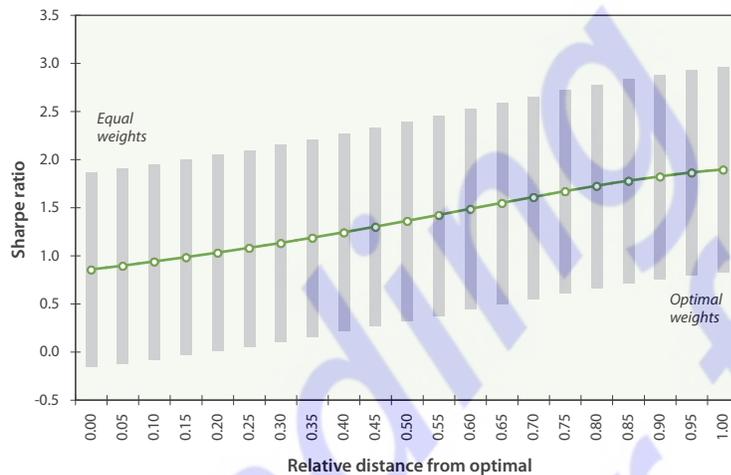
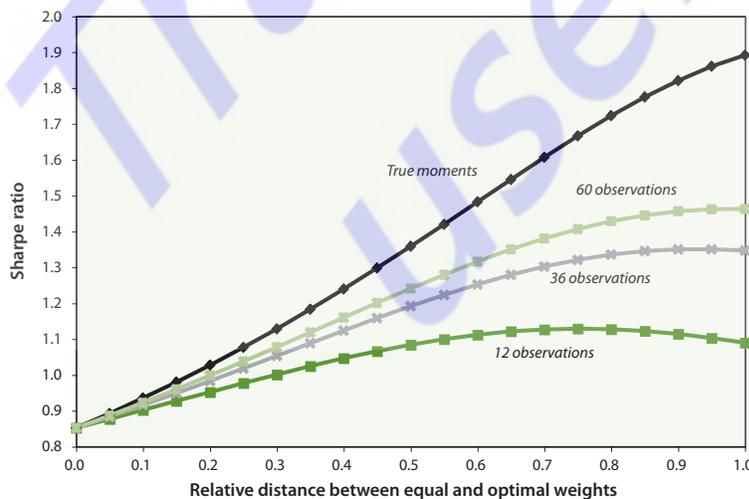


Exhibit 8
Sharpe ratio as a function of distance from optimal weights (versus number of observations used to estimate moments of return distributions)



estimates around the true mean is very wide. In contrast, when estimating the implied correlation of a group of assets, the distribution of the estimate can actually be smaller than it is for more highly correlated assets. As shown in Exhibit 6, the standard deviation of the implied correlation estimate is actually smaller for six assets when correlation is zero than when it is higher. This then explains the shape of the (American) football scatter one sees in Exhibit 3.

The work we did for *There are known unknowns* (Newedge, Sept. 10, 2007) convinced us of the need for real humility when working with real portfolios. For one thing, the distributions of Sharpe ratios for reasonably short holding periods can be very wide. Exhibit 7 shows a 1-standard-deviation band around the Sharpe ratios one might expect from volatility weighted portfolios (at the left) and fully optimized portfolios (on the right). The lesson we took from this exhibit is that while there are things we can do to improve a portfolio's expected performance, we cannot seriously expect to detect the results of our work in track records as short as a year.

Furthermore, we are hampered in any effort we make to optimize portfolios by the fact that we are always working with estimates of the parameters of return distributions rather than with the truth. Exhibit 8 provides some insight into just how much of a drag imperfect knowledge places on what we hope to do. This exhibit was the result of assuming that what we had estimated for these assets was in fact the truth, using those values to produce several thousand return paths that were consistent with those values, and then estimating over look-back periods ranging from 12 to 60 months. Given the way we've constructed the experiment, the estimates using 12 months would have the biggest errors. The curves drawn for 36 and 60 observations show how the expected results improve as the quality of the estimates improve.

The apparent drag of having to work with

imperfect knowledge is huge. Moreover, if the estimation errors are large enough, the so-called “optimized” portfolio might not even be better than a compromise solution. With a 12-month look-back, for example, the highest expected Sharpe ratio comes from a portfolio that is about halfway between a simple volatility weighted portfolio and a portfolio optimized using estimated returns, volatilities, and correlations.

Index methodology and construction

With this research in mind, the Short-Term Traders Index was constructed with three things in mind. It had to be simple. It had to be robust. And it had to reflect a minimum of curve fitting. To be included, a manager’s trades must have an average holding period of 10 days or less. Further, the manager must trade in two or more broad market sectors – equities, fixed income, foreign currency, and commodities – be open to investment, and be willing to report daily returns.

In its current form, 90 percent of the weight is given to managers who have \$100 million or more under management at the time the index is reconstituted. These managers’ returns are equally volatility weighted – giving higher weight to low volatility managers and lower weight to high volatility managers – using estimates based on the previous year’s daily returns. These weights are subject to

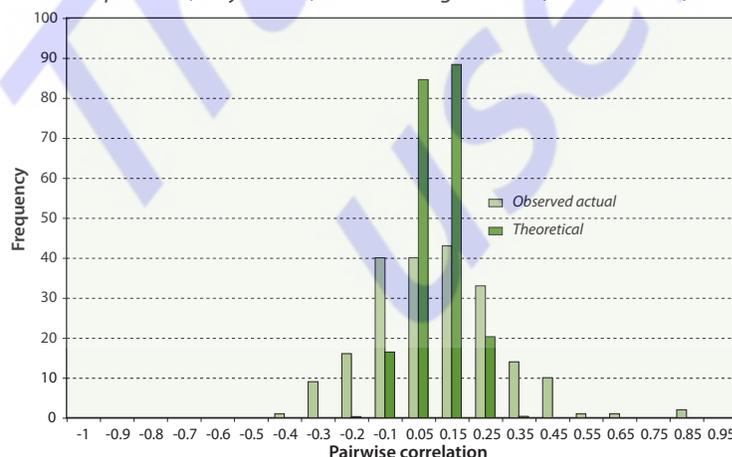
minimum and maximum constraints to avoid serious distortions from errors in volatility estimates. At this writing, this part of the index comprises 14 programs whose combined assets under management equal \$6.9 billion.

The remaining 10 percent of the index is assigned at the Index Committee’s discretion to managers who are representative of the short-term space and meet all criteria except the \$100 million minimum. At this writing, this part of the index comprises seven programs whose returns are equally weighted.

Exhibit 9
Alternative Edge Short-Term Traders Index



Exhibit 10
Actual and theoretical distributions of pairwise correlations among 21 STTI components (daily returns, 6/1/07 through 5/30/08, mean = 0.104)



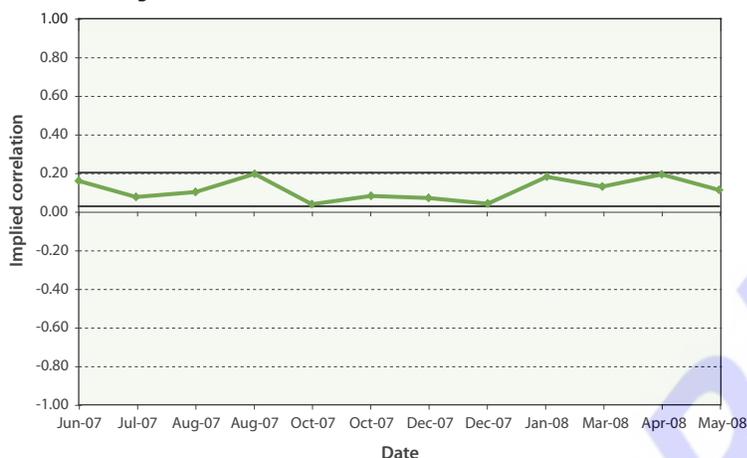
How low are the correlations?

Going into this exercise, we expected the correlations of returns among the managers to be low, and we expected the correlations of their returns with those for longer-term CTAs and for other investments to be low. The exhibits in this section bear this out.

Pairwise correlations

Exhibit 10 shows the distribution of pairwise correlations for the 21 programs, estimated using daily returns, for the period June 2007 through May 2008. The average of these correlations was 0.104. To put this distribution in perspective, we have overlaid the theoretical distribution of estimates one would get if the true pairwise correlation for all managers were

Exhibit 11
Implied correlation for 21 STTI components
 (Average = 0.116)



in fact 0.104. To us, the interesting lesson in comparing the two is that the observed distribution is wider than it would be if all of the CTAs actually had the same low correlation with one another. At the same time, it seems that one could construct the observed distribution, with the exception of the 0.85 outlier, from a blend of distributions with a fairly small range of true (and low) correlations.

Time series of implied correlations

While Exhibit 10 used the entire year of returns to estimate pairwise correlations, Exhibit 11 shows how the average or implied correlation estimates varied over the year. Each point is the average pairwise correlation estimated

using that month's daily returns. Overlaid on this series are two lines that represent a 2-standard-deviation band around the average of 0.116. This band was reckoned assuming that the true pairwise correlation for all managers was the sample mean. In reality, the actual 2-standard-deviation band is wider than this because some of the true correlations are less than the mean used here, while some are greater. In any case, a time series like this is a useful diagnostic for monitoring the overall correlation among managers in the index.

Correlations with returns of other investments

Exhibit 12 reports the correlation of daily returns with those on other investment vehicles including the Newedge CTA index, the S&P 500, the Lehman Composite Bond Index, and the HFR Equity Hedge index. As one would expect, the correlations of STTI returns with those of other investment vehicles are all fairly low – some slightly positive, others slightly negative. Even its correlation with the Newedge CTA index, which comprises the 20 largest CTAs, is only 0.235.

Exhibit 12
 From Jun 07 to May 08

	STTI	Newedge CTA Index	S&P 500 Index	Lehman Composite Bond Index	HFR Equity Hedge Index
STTI	1.000				
Newedge CTA	0.235	1.000			
S&P 500	-0.246	0.031	1.000		
Lehman Composite Bond Index	0.104	0.068	-0.395	1.000	
HFR Equity Hedge	-0.131	0.299	0.771	-0.365	1.000

Correlations of returns calculated over longer periods

Correlation estimates can be sensitive to the choice of period over which returns are calculated. For example, the use of daily data may introduce noise because of different settlement times throughout

the day. To see whether the appearance of low correlation might be a result of using daily data, we calculated correlations of STTI and Newedge CTA returns for non-overlapping periods of five, 10, and 20 days, corresponding roughly to weekly, biweekly, and monthly returns. Exhibit 13 shows that nothing much changes as we lengthen the period used to calculate returns. The correlations appear to increase slightly as

Exhibit 13

Interval in days	Count	Correlation	Stdev
1	264	0.233	0.058
5	52	0.242	0.131
10	26	0.310	0.179
20	13	-0.083	0.276

we go from one to five to 10 days, but nothing that is inconsistent with sampling error. And then, when we go to 20 days, the estimated correlation is actually slightly negative. But again, the change is well within the range of sampling error one would expect.

Why are the correlations low?

The most obvious place to look for an explanation for the low correlations we observe is the ultimate source of all returns – the markets to which all traders have access. Two traders’ returns at any moment and during any particular period can be written as:

$$R_{it} = \sum_{m=1}^M w_{imt} R_{mt}$$

$$R_{jt} = \sum_{m=1}^M w_{jmt} R_{mt}$$

where i and j are different traders taking positions in any or all of M markets. Since both traders have access to the same markets, their returns can be different only if the weights they assign to each of the markets are different, either in size or sign.

What we suspect, of course, is that short-term traders’ trades overlap very little with those of longer-term traders. Either they are not in the market at the same time or the directions they take in the markets may be independent of the directions taken by longer-term traders.

Holding period and return correlation

A simple experiment bore this out. Using five years of real-time data for S&P 500 futures from 2003 through 2007, we compared the results of a simple M-day/N-day breakout model implemented for various look-back periods. In particular, we implemented the rule as follows:

Entrance:	If price > M-day high, go long; If price < M-day low, go short
Exit:	If price < N-day high, exit long; If price > N-day low, exit short

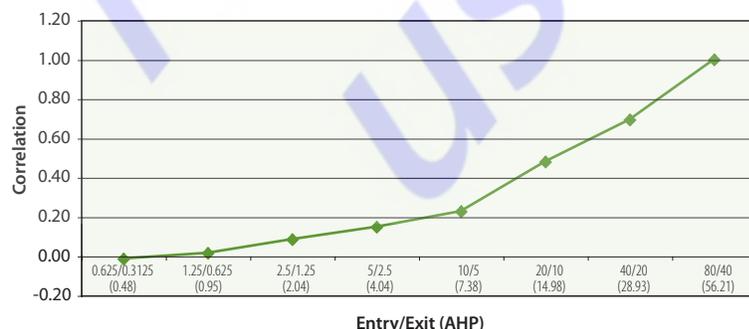
for eight look-back periods: 80/40, 40/20 and so on, reducing the times by half until we got to 0.625/0.3125. In implementing these rules, we assumed that once a trade is initiated, a new position cannot be taken until the older trade is closed out.

The effect of shortening the breakout period can be seen in Exhibit 14, which shows the correlation of each strategy’s returns with those generated by the 80/40 breakout rule. (The resulting

average holding period for each is shown in parentheses.) We see that the correlation falls as we shorten the holding period until, by the time the average holding period is under 10 days, the correlation is in the neighborhood of 0.20. By the time the holding period is less than a day, the correlation with 80/40 returns is about zero.

What we see in Exhibit 14 is consistent with a sense we have that there are far more

Exhibit 14
Correlations with 80/40 breakout model



ways to slice up a time series if our holding period is short than if our horizons are long. This is borne out in Exhibit 15, which shows how many S&P 500 trades each rule generated during the five-year period. The 80/40 rule produced a total of eight trades with an average holding period of 56 days. The 0.625/0.3125 rule, in contrast, produced almost 1,400 trades with an average holding period of less than half a day.

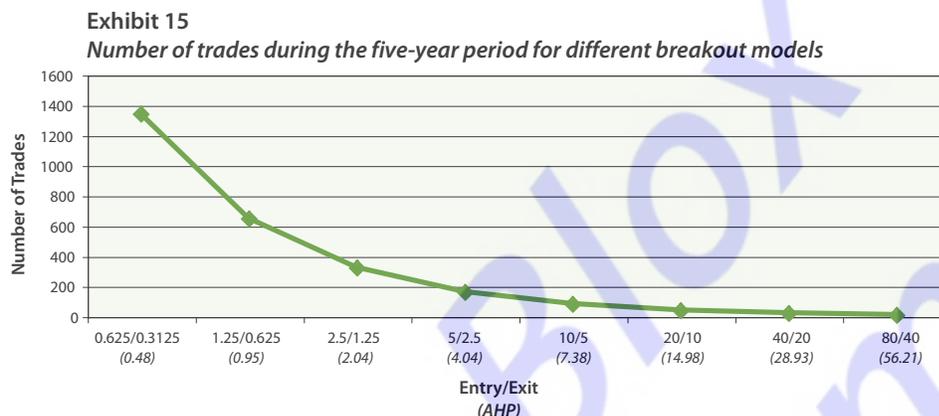


Exhibit 16
Fraction of the time in the market with the same position as the 80/40 breakout rule

Lookback	Fraction
80/40	53.8
40/20	42.0
20/10	28.9
10/5	22.3
5/2.5	23.4
2.5/1.25	21.3
1.25/0.625	19.1
0.625/0.3125	19.8

As a result, the amount of time that traders spend in the market at the same time and in the same direction falls as the holding period is reduced. Exhibit 16 bears this out by showing the fraction of the time each rule is both in the market and in the same direction as the 80/40 rule. The first number indicates that the 80/40 rule was in the market 53.8 percent of the time. As it was, each rule was in the market roughly the same percent of the time as the 80/40 rule, but the overlaps tended to shrink as the holding period was reduced. The shortest holding periods were in the market and in the same direction as the 80/40 rule only about 20 percent of the time.

More evidence on correlation

Pierre Villeneuve, managing director of Mapleridge, offered further insights into why short-term strategies exhibit such low correlation. In a research presentation at our research forum in Bordeaux, he compared short-term and long-term trading strategies that were purely random on the one hand and that had trend following features on the other. He also isolated the profitable strategies in both cases to see how focusing on “successful” trading programs might affect the results.

His purely random model was described as:

Start with a random position long/short (+1/-1)

Each day change position with probability p

In his examples:

- Short-term p = 0.5
- Long-term p = 0.008

So the average holding period for the short-term model was 2 days [= 1 / 0.5], while the average holding period for the long-term model was 125 days [= 1 / 0.008].

His trend following model was described as:

Start with a random position long/short (+1/-1)

Each day if:

50-day moving average of returns is against current position, then change position with probability p (more likely to change position)

50-day moving average of returns is with current position, then change position with probability p^2 (less likely to change)

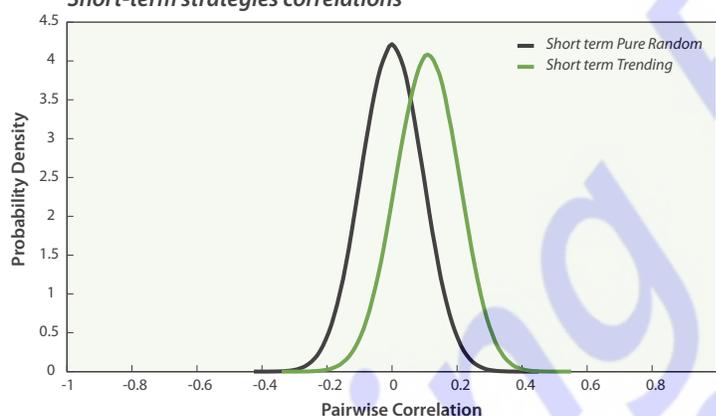
In his example:

Short-term $p = 0.5$

Long-term $p = 0.02$

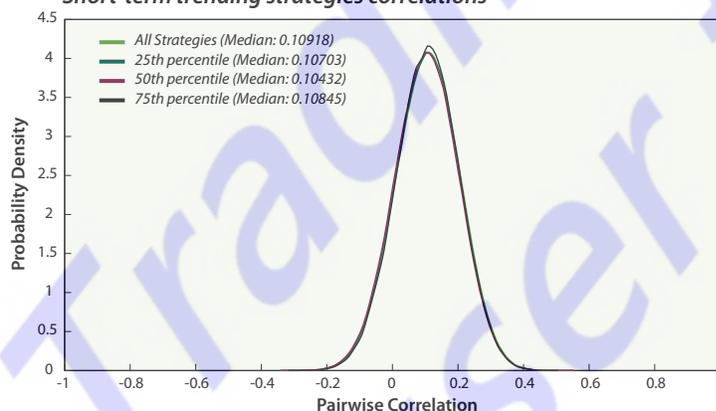
In both cases, and for both short-term and long-term models, he generated 2,000 trials using a continuous S&P 500 futures series from 2005 through 2008.

Exhibit 17
Short-term strategies correlations



The distributions of pairwise correlations for the random and trend following short-term strategies are compared in Exhibit 17. It is apparent that the introduction of a trend following component to the model increased the average correlation from 0.0 to about 0.1. In Exhibit 18, however, we find that the distributions of pairwise correlations for the trend following models were independent of whether they made money or not. The distribution of correlations for those runs that produced gains in the top 25 percent was no different than the distribution of correlations for returns from all runs.

Exhibit 18
Short-term trending strategies correlations



In Exhibit 19, we find a similar result, although the effect of introducing a trend following component to the long-term strategies was more pronounced. The average pairwise correlation increased from 0.0 for the long-term random strategies to 0.34 for the long-term trend following strategies. Further, in Exhibit 20, we see that isolating the profitable trend following strategies – represented here by the top 25 percent of the distribution of returns – has an additional effect on the distribution of correlations. The average correlation for the 75th percentile (the top 25 percent) was 0.59.

His explanation for the differences in the short-term and long-term results agreed with our suggestion that the possible overlaps for short-term strategies are much fewer and much smaller than what one finds with long-term strategies. There are simply a lot more ways to trade markets in the short-term than in the long-term. As a result, even the successful short-term outcomes tend not to overlap with one another. With long-term strategies, however, there are comparatively few profitable trades, which suggests that overlapping positions are highly likely.

Exhibit 19
Long-term strategies correlations

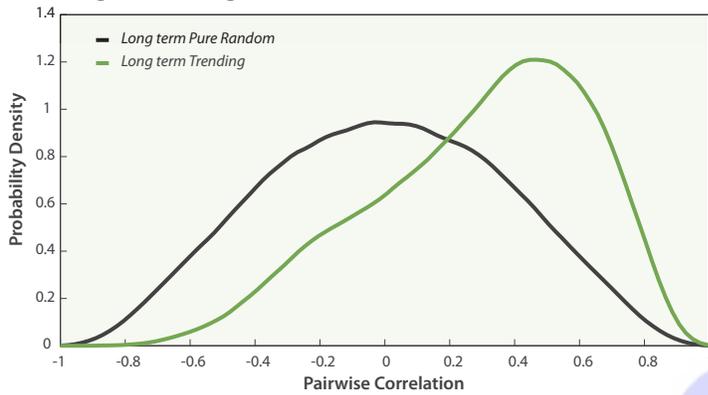


Exhibit 20
Long-term trending strategies correlations

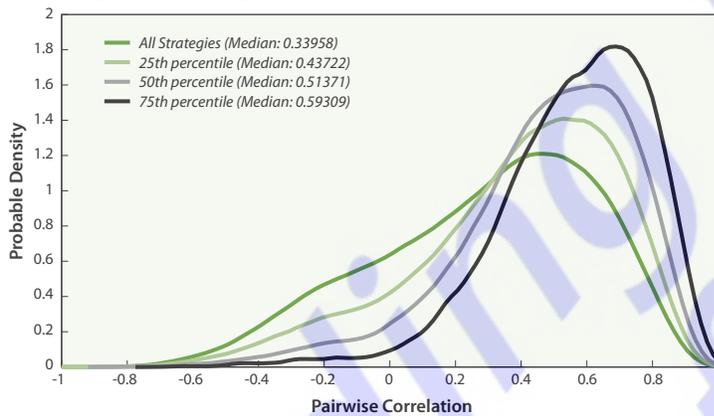
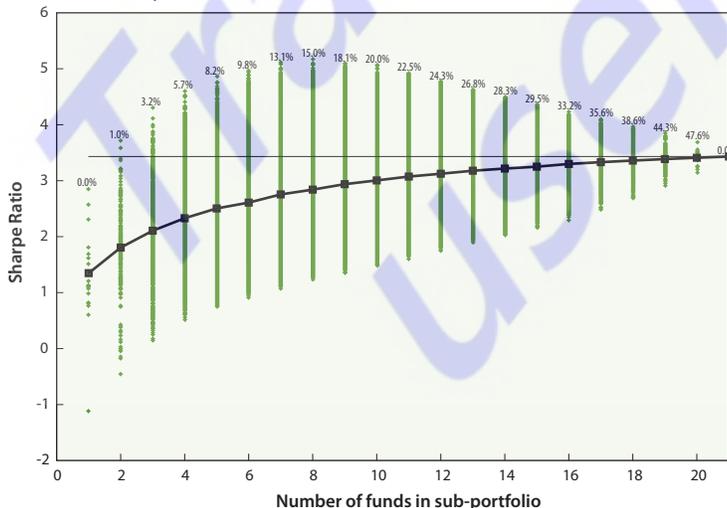


Exhibit 21
Sharpe Ratios for all possible STTI constituents sub-portfolio (2007.1 to present)



Why are there not more short-term traders?

If short-term trading has such desirable properties for portfolio management, why have more traders not been attracted to the space? The most obvious answer seems to be that the barriers to entry are high. For one thing, results are highly sensitive to transactions costs and market impacts. As a result, the amount of care and research required to ferret out profitable strategies and to keep transaction costs at an absolute minimum tend to be much greater than for longer-term trading strategies.

Although the number of possible strategies is very large, the capacity of any given strategy is questionable. At the same time, the electronic revolution in futures trading has been hugely effective in reducing transaction costs. Liquidity has improved immensely in many markets and the drag of market impact has been brought down substantially. As a result, strategies that might not have been profitable in the past may well be profitable now.

Replicating the index

Investors attracted to the idea of investing in this space may find it either impractical or impossible to invest in all 21 programs in the index. While all 21 are open for investment, not all take managed accounts. This increases the cost of investing because the leads and lags that are part of investing in funds can be considerable. Also, investors may have internal constraints that prevent them from investing in one or more of the managers, or find it too costly to conduct due diligence for all 21 programs.

To get a sense of what one might expect from investing in subsets of the 21, we considered all possible portfolios ranging from 1 program (single program) to 20 programs and calculated the Sharpe ratios one would have realized over the period from January 2007 through May 2008. The results of this work are shown in Exhibit 21.

For each portfolio size, the exhibit shows several things – the range of possible Sharpe ratios,

the average Sharpe ratio, and the fraction of outcomes that produced a Sharpe ratio greater than the 21-program portfolio. The chief lessons are those we know from the rules that govern diversification and risk reduction. Most of the benefits from diversification are reaped by the time one has reached 10 managers. Even so, the likelihood of doing as well – at least by this standard – is only about 20 percent.

It is also apparent that investing in three or four short-term traders does not capture the full benefits of investing in this space. Only 3 percent of the 3-program portfolios outperformed the index.

Red flags, cautions, and managing the index

We should note that we rely to a large extent on self-reporting and market discipline to keep the index on track. We do not audit the managers' trades to make sure that their average holding period is 10 days or less. Rather, we rely on managers telling us what they tell their clients. Also, the \$100 million criterion for assets under management for 90 percent of the index gives us confidence that these managers have passed muster with substantial investors.

Thus, our greatest concern in managing the index is with correlation. Given our prior beliefs about average holding period and return correlations, we are satisfied with the correlative properties that the index and its components have exhibited so far. For us, a red flag would be a change in implied correlation that was inconsistent with what one would expect. As we showed in Exhibit 11, a 2-standard-deviation range around a true correlation of 0.116 would be from 0.03 to 0.20. Because we think the actual ranges of true pairwise correlations occupy a range around 0.116, our tolerances would be wider than this.

Conclusion

The development of the AlternativeEdge Short-Term Traders Index has been a constructive exercise that puts to use insights from the research we reported in *Superstars versus teamwork* and in *There are known unknowns*. Diversification when correlations are low produces superior portfolios. Moreover, when working with groups of 10 or more assets or managers, correlations tend to be predictable and we actually have more confidence in low correlation estimates when correlations are low. In the case of short-term traders, we have every reason to believe that their returns will exhibit low correlations with one another and with those of longer-term traders because of the richer set of trading opportunities they face. So far, our expectations of low correlation have been borne out. Moreover, we are confident that we have done a good job of constructing a collection of managers who are highly and broadly representative of the short-term trading space.

Acknowledgements

Our various research forums have produced many valuable insights for us, and this work reflects some of them. We want to acknowledge Rishi Narang of Telesis Capital for encouraging us to drill down on why the shape of the correlation scatter in Exhibit 3 was what it was. This led us to the realization that correlation estimates for groups of low correlation assets can actually be better than those for higher correlation assets. We definitely want to thank Roy Niederhoffer of R.G. Niederhoffer Capital Management for making the connection between low correlation and holding period and for suggesting that we pursue the construction of a short-term trader index. And we want to recognize Steve Evans of Tudor Investment Corporation and Rafael Molinero of Molinero Capital Management who cautioned us at length about the possibility of obscuring true correlation relationships because of an inappropriate choice of the period over which one calculates returns.

This page intentionally left blank

Trading Blox
user forum

AlternativeEdge series

- **How important are daily return data?**
- **Understanding drawdowns**
- **Every drought ends in a rainstorm**
- **Measuring market impact and liquidity**
- **Superstars versus teamwork**
- **What you should expect from trend following**
- **There are known unknowns**
- **Costs of active management**
- **Correlations and holding periods:**
The research basis for the AlternativeEdge Short-Term Traders Index

Please feel free to contact Galen Burghardt (312-762-1140, galen.burghardt@newedgegroup.com) or Ryan Duncan (312-762-1122, ryan.duncan@newedgegroup.com) with questions about the AlternativeEdge series of research notes.