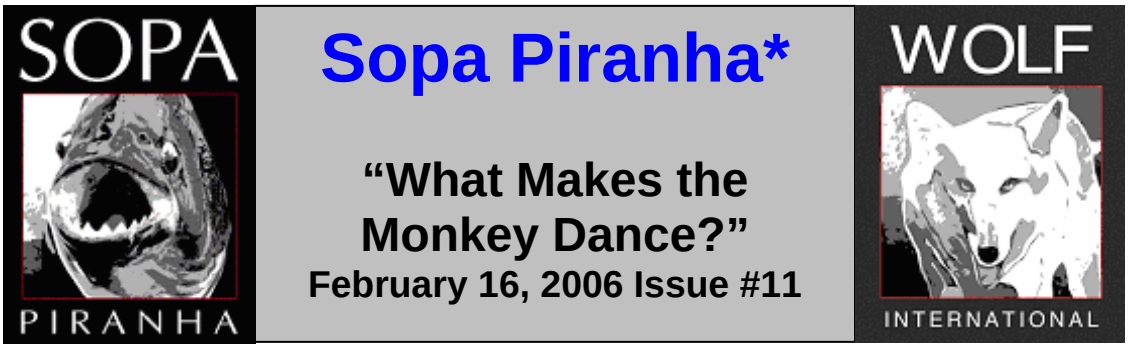

Quarterly Commentary: This quarter's topic is **systematic trading** with special emphasis on statistically based market neutral. Our portfolios have been dominated by growth themes for 3 years and it is time for us to create some more balance with market neutral systems. As I have often stated, a well balanced high return market neutral portfolio can be leveraged to earn 15%. It also provides diversification that allows for more risk in other parts of the portfolio. We should all remember that some of the world's best track records in hedge funds are built on this concept.

Last quarter, I focused on the short side of the market noting that short side attribution is improving. The **attribution is also improving** in systematic trading. The study below includes information about which **kinds of systems** work and **best practices** in the market. I have included comments on the market **environment that is best** for the strategy. As usual I start simple and move to the complex.

Section One: What is statistical arbitrage? Over the long term, stat arb managers as a group under perform hedge fund indices. This is entirely due to lack of beta. A strong statistical arbitrage program laid over pure 100% beta will outperform the hedge fund index. Therefore stat arb is one of the best remaining strategies for hedge funds who want portable alpha. Of the 70 hedge funds that we located who advertised some form of statistical arbitrage, **only 40** were actually doing statistical arbitrage. Like many terms in the hedge fund industry, statistical arbitrage has been bastardized to mean any equity trading with equal weightings using some form of statistical analysis. Most of what we found was glorified pair trading. Pair trading is long one stock and short another within a sector. We also uncovered the fact that many of the stat arb funds in the available indices are not stat arb. This apparent lack of a clear understanding of real statistical arbitrage suggests I need to define in the broadest way what a sophisticated stat arb effort looks like. Simply stated a statistical arbitrage program is a **bottom up**, **quantitative**, **beta neutral equity** strategy where statistical analysis and development drives every part of the investment process. Complex historical **patterns** are derived from various **time series** to identify systematic relative value opportunities.

I. Statistical arbitrage for beginners

- **History:** Statistical arbitrage began in the mid 1980's. Morgan Stanley is often given the credit for bringing it into the mainstream. Relatively simple quantitative models were an outgrowth of the successful derivative business. The strategy grew rapidly through the 90's and most large investment houses entered the game. Trading margins were wide and so the quality of the model was less important. Data and computing power were both expensive. By the late 90's volatility was extremely high and costs were declining. **Stat arb funds proliferated until 2004 when volatility collapsed**. Recently a consolidation has occurred and only the most sophisticated shops are thriving.

- **Viable markets:** The only viable markets are those that permit active shorting and are well regulated. Over 50% of all statistical arbitrage is executed in the US listed equity market. This is because frequent trading requires ample liquidity. Less liquid markets are less efficient and provide more opportunity. This opportunity can be difficult to capture consistently because of higher market impact and even the lack of good historical data. **Therefore stat arb managers are constantly engaging in a tradeoff between a desire for liquidity and the need for inefficiency.** Larger stat arbs are more likely to operate in many markets and their large capital base makes execution difficult. Smaller stat arb managers may not have the experience or resources to operate in many markets. Viable markets include the US, much of Europe, Japan, and Australia. More recently South Korea and Singapore are getting attention.

- **Simple signal generation:** A statistical arbitrage system generates a trade signal when the portfolio is **optimized with new data**. A very simple example of a forecast signal is a mean reversion model. The system computes the mean return of all stocks within an industry and those that are far enough above the mean are sold while those that are far enough below are bought. When the stocks revert to the mean return of the group, the trade is reversed. Multiple signals are used to increase total return forecasting power. These are used by the portfolio optimizer to generate orders.

- **Simple trade entry:** Orders are done either through an in-house execution system or they are given in batch to an electronic broker. The orders are often split into smaller batches and fed into the market in a size that will minimize price impact. High frequency strategies will normally have execution built into their portfolio construction. This means the **market response to a trade** is being monitored tick by tick for its impact on the signal generation.

II. **Data on the statistical arbitrage industry:** Studying historical data on stat arb managers in batch is nearly impossible. Systems vary widely by market and model type. They have also migrated considerably even in the last three years. The market environment has changed dramatically. The factors that most affect stat arb have been among the **most volatile market measures recently**. Specifically volatility on equities has dropped nearly 200% while dispersion has dropped over 100%. This forced us to analyze the statistical arbitrage managers in groups with similar models and style drift. Even then, we can get no more than three years of good comparison. This led to a need to create our own data from scratch and to analyze what is possible.

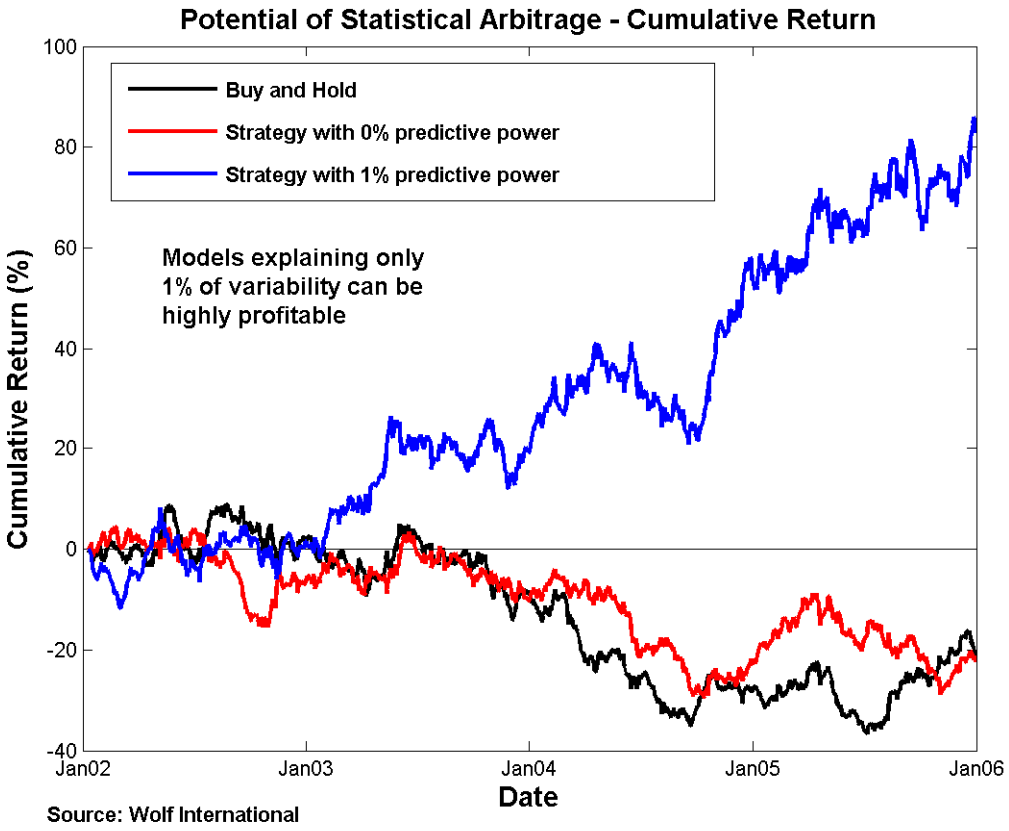

The chart above shows graphically a fund we created. There are three lines. The black line is random market data, the red line has no predictive power, and the blue line is our fund with 1% predictive power. It shows with simple 1% predictive power, a system can produce huge returns due to repetitive turnover. The black line on the chart above shows the cumulative return of a hypothetical instrument that has completely random daily returns with mean 0% and standard deviation 1% (**normal distribution**). It is just one example of unlimited (random) possibilities. We

call this the Buy-And-Hold strategy. We then created a trading strategy that uses a **single forecasting signal**. This signal is constructed to have pre-set forecasting power to predict the returns of the black line. In other words, the signal that's created is also drawn from a normal distribution with mean zero and standard deviation of 1. But the "signal" distribution is correlated to the "instrument" distribution at a defined level (using a fabricated covariance matrix between the distributions). This signal is then used to trigger buy/sell actions in a very straightforward way. Transaction **costs are assumed to be zero**. The red line shows the cumulative returns of a strategy which uses a signal that has an R-squared to the returns of the instrument of 0%. In other words, the signal is completely uncorrelated. Buy and sells are random and the strategy should (on average) not do better or worse than the Buy-and-Hold strategy when transaction costs are zero.

The blue line is a strategy that uses a signal that explains just **1% of the variability** of the underlying instrument and gives you a good idea of the **potential of stat arb** (R-squared is 0.01 and correlation of signal to instrument returns is 10%). It shows that signals only need a limited amount of forecasting power in order to be highly profitable. However, it must be mentioned that achieving even small amounts of **consistent predictability** is not trivial, nor is the process of capitalizing on such signals **while limiting alpha destruction**.

It is also critical to understand that we could have created a strategy with no forecasting power, similar to the red line, which outperforms the predictive strategy (blue line) over the same time period. The chances of this happening at random are small but such a fund would seem **to be robust** on the surface when analyzed using standard industry measures. However, only a precise understanding of the **best practices** of stat arb taken in conjunction with the predictive models would reveal that we created a complete garbage fund. Statistically, the upward return stream would be unsustainable. The random positive returns would inevitably revert to zero. Such a fund ignores the essential components of stat arb that must be closely scrutinized in conjunction with the predictive models. This leads us to best practices in statistical arbitrage.

Section Two: What are best practices of a good statistical arbitrage fund? Five stages of a strong stat arb are: **I) universe construction, II) signal generation, III) portfolio construction, IV) automated execution and V) risk management.** All are a source of alpha and all are important.

I. **How should universe construction be done?** A Statistical arbitrage system will include any stock that meets its filtration criteria. All other factors being equal, an increase in the number of names results in a more diversified portfolio and thus allows for more **leverage**. Some key points in universe construction are:

- **History of data used in model creation;** The length of time depends on the forecast horizon of the signals. This is particularly important if new data sources are made available to the market because these are often a big source of stat arb alpha. Can be incorporated quicker with short term models that need less historic data to prove profitability.

- **Frequency of data update;** The best systems continually update their universe of stocks. Since most stat arbs use somewhat similar filtering criteria, adding new names quickly gives an edge to the manager.

- **Liquidity;** In today's fragmented market, it is important to only include *accessible* liquidity. A stock may be reasonably liquid, but if trading volume is fragmented across markets, the manager may not be able to tap into all liquidity sources. A minimum volume such as **100 thousand shares per day traded** is helpful.

- **Market Capitalization;** A manager should include stocks that have a threshold market capitalization. This helps avoid small companies that are too volatile and hard to model. It can help **avoid takeover targets**.

- **Stock Price;** The stat arb should only trade stocks that have a minimum price such as **5 dollars per share**. This helps avoid delisting risk and a high ratio of **minimum price change to share price**.

- **Clustering;** A stat arb could run one trading universe as a whole and deem each stock to have similar characteristics. However, it is quite typical to sub divide the portfolio in separate clusters that are made up of stocks that are related. Some stocks may be dropped if they are part of a cluster that has too few names. A cluster may be determined by a data provider such as Barra but there is an advantage to **proprietary analysis of clusters** based on statistical methods.

- **Other methods;** Other more sophisticated factors are often proprietary and I don't want to give away the secrets of our managers. It is important to know that without an edge in universe construction, alpha will be less and it will be difficult to compete. Some people say that it is impossible to compete in the stat arb space without a massive infrastructure and huge amounts of data. **This is incorrect**. There are

highly successful stat arbs that only use daily price data to create their forecast signals. **One liquid market is sufficient** to run a successful stat arb strategy and a US centered product is normally launched first.

Stat arb profitability comes from consistent alpha harvesting across *all* stocks in its universe. It is very unusual to have a single stock account for a large portion of profits over a single year. This is one sign of a poorly constructed stat arb system. As a result, stat arbs benefit greatly from the **largest possible universe**. Stat arbs with small and static universes normally focus on a small area of predictability that might do quite well at low ramp level, but can **fall apart dramatically when assets increase**.

Stat arbs are the most sophisticated users of market data and they use a lot of it. Any data that can be represented by a time series and that has some influence on market prices will be utilized to enhance the models and increase the forecasting power of their signals. All stat arbs use **daily price and volume data, but could also include tick data, limit order book data, balance sheets, analyst estimates and short interest data**, just to name a few. A short term strategy may use tick and book data, but might ignore debt to equity ratios since they don't provide short term predictability. Longer term strategies may do the opposite. It all depends on the timescale of the signal.

II. How should signal generation be done?

There are about **25 different sub strategies** in statistical arbitrage. They fall into four broad groups; **mean reversion, psychological extremes, short term momentum, and corporate quality**. Any one of the four can produce an excellent fund especially if the group is fleshed out with all of the possible types of signal generation it can have. The principal advantage of using more than one group is not return, it is leverage. The diversification benefit can be used to increase leverage quite a lot. This is why shops with many models tend to have **higher Sharpe ratios** but not necessarily better absolute performance in any one year.

Signal creation is the core of every stat arb and is always based on detecting consistent **repeatable** patterns in historic data. Stat arb designers look for high probability of frequent repetition so that each trade produces a relatively low rate of return. Creation of these signals can be accomplished with a variety of mathematical methods. Examples of methods include PCA (principal component analysis), co-integration and auto regression. All methods can be used independently **or in combination** to generate stable non-correlated alpha.

No matter what the mathematical method, the objective is to find factors or deviations that predict price movement. Often the statistical method is designed to **normalize the data so that some factors are neutralized** in order to isolate others that are a predictive signal. Other models derive signals from market or cluster hedged returns. Forecasts are looking for excess returns over market returns and are thus **market neutral by design**. If a system has very high beta or extreme variable beta, it should be re-categorized as a trend follower even if statistics are driving the model.

The forecast signal is tied to the **forecast horizon** or time period. One model type will often be used over more than one time period. One time period may have more than one model predicting price. But in general terms, there is a specific time that makes economic sense. An earnings model is **timed to match** the earnings release season. The most common time frames for full portfolio turnover are about 5 days, 20 days, 65 days, and 130 days. It is not an accident that these coincide with the number of business days in a week, month, quarter and half of a year.

The historical period that is used for finding patterns is called “in sample”. A bigger pool of signals to predict future returns is not always better. Only if signals **complement each other**, are fairly uncorrelated, and are consistent performers should they be included. It is always possible to add mediocre signals that improve forecasting power during in sample periods only to see a big drop in performance when used on historical data that was *not* used in deriving the signals (“out of sample”). This means the additional signals were explaining noise of the in sample period, which was not present in the out of sample period. This ***over fitting* of the data** is one of the biggest challenges of stat arb and needs to be closely monitored. On the other hand, having alpha concentrated in just a few factors is risky when the performance of those signals inevitably decays.

Once signals are chosen, the independent signals are then combined in a fitting procedure to determine their **optimal weights** in order to maximize explanatory power over the in sample data period. If a stat arb manager wants to go to a new market, he must start over. Concepts can be migrated to other markets, but the detailed methods cannot normally be used.

III. **How should portfolio construction be done?** Statistical arbitrage can be thought of as a process of identifying alpha and then working to enhance it through leverage while not destroying it through execution and risk management. Portfolio construction is the step that can dramatically enhance the ability to leverage. Some managers run a

1x1 system and do well. Others are running **20 systems with overall leverage of 6x6**. Portfolio construction is often done with numerical optimization that uses an iterative process to solve for the portfolio that maximizes risk adjusted return. At the same time the optimizer is limiting the portfolio with constraints that can change. Typical inputs to this process are as follows:

- **Stock specific forecasts of future returns**. This is derived from the process described above.
- **Transaction cost model.** It is critical to have a reasonable way of limiting market impact. With similar return forecasts, the transaction cost model will allow **bigger trades in high liquidity stocks** while limiting those for low liquidity stocks. On the other hand, it may penalize trading of a highly liquid stock if the forecast is too small.
- **Current portfolio positions.** These are needed to estimate the transaction **cost incurred** when trading away from current holdings.
- **Risk model.** Managers estimate risk by using a proprietary or third party fundamental or numerical **factor model**. The risk model is not just for risk management. It feeds back into the predictive model in the most sophisticated applications.
- **Constraints on stock, cluster, sector, market, factor, dollar and volatility exposure.** These define the boundaries in which the optimizer can construct the optimal portfolio. The best stat arbs are normally **cluster, sector, market, factor and dollar neutral**. If these things are kept near neutral, a lot of leverage is possible while still limiting exposure to large left tail events. This is particularly true when volatility is low. If volatility is increasing, good systems automatically reduce exposure to certain stocks in order to satisfy the total risk constraint.

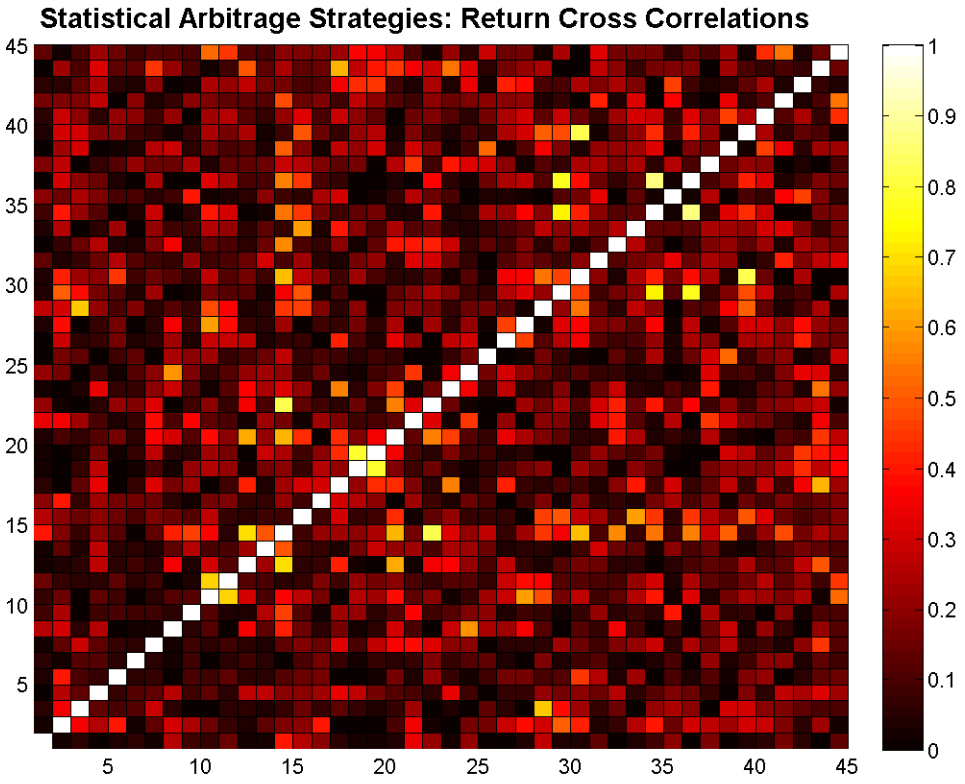

Source: hedgefund.net / Wolf International

Portfolio construction is beneficial to reduce risk but it should be noted that statistical arbitrage is by its nature not even correlated to itself. The heat map above shows the correlation among different statistical arbitrage funds. Stat arb is not correlated to other hedge fund strategies and as the map above shows, it is **not even correlated to other stat arb models**. This low correlation is what makes stat arb worth doing.

Stat arb is also a remarkably beneficial part of a hedge fund portfolio. It is one of the very few strategies that likes rising volatility. When combined with strategies that are effectively short volatility, such as credit, it really makes a portfolio sing. Some of the world's best track records prove this point. For those who believe stat arb is a dead strategy, you might want to call five of the world's best hedge funds and ask them why they spend so much money on it. Top managers are still posting impressive numbers.

IV. **How should execution be done?** Execution is another often overlooked source of hidden alpha or said another way, if it is done properly, free alpha is not destroyed. It is crucial for stat arbs since they turn over their portfolios often and create many small orders with small gains per trade. This is most important for the high frequency stat arbs where transaction costs are a large fraction of their profits. There is **no longer a need to develop execution systems** in house, since innovation and competition have provided the market with a wealth of advanced execution systems. Still, an in-house execution system can give a slight advantage for high frequency stat arbs. Some locate their trading system as close to the electronic matching engines as possible. Monitoring of realized versus modeled transaction costs is critical for the portfolio optimizer to work properly.

Some stat arbs have taken the novel approach of **integrating the portfolio construction** module with the execution system. In other words, instead of a single set of orders for the whole portfolio that the execution system is to fill during the day, the optimizer is actually constantly updating the portfolio during the day. This is not only important for high frequency trading strategies. The forecasting power of some medium term models is concentrated in the early part of the forecasting period. The execution system may not be aware of this and execute in a volume weighted average pricing fashion, missing some of the alpha that the model properly predicted.

V. **How should risk be controlled?** Advanced risk control should be **included in the design** and development of each component. From universe selection and signal creation, to portfolio construction and execution, each piece has specific risk control challenges that need to be addressed. Most stat arbs use fundamental or numerical factor models that describe the exposure of each stock in the portfolio to these factors. The residual risk that cannot be described by these models is called **stock specific risk and should ideally be the only source of risk** in a pure market, sector and factor neutral stat arb.

Only in very specific situations should the stat arb model be overridden. It is almost always detrimental to intervene with the systematic order flow on a consistent basis, since it violates all the statistical analyses that have been performed in building a well-conditioned model. Intervention is only warranted when specific **uncommon events occur** that, by their infrequent nature, are not well explained by the forecast models. A typical example is a merger where a stat arb holding in the acquired company would normally be liquidated. Other examples would be a dramatic **drop in liquidity** below preset criteria, hard to borrow events or regulatory changes that

force stock liquidation. Any stat arb manager that interferes with his own order flow consistently should not be considered a true stat arb.

Summary of section two: Statistical arbitrage should follow a five stage process that is fully integrated at every level. Once a system is fully designed, very few people are needed and the strategy is somewhat portable. This is why managers are so careful not to reveal their market edge. When evaluating the strategy some of the pitfalls to avoid are as follows:

Not systematic
Small or illiquid universe (e.g., only focus on small caps)
No or improper use of transaction cost model
Signal set has not been enhanced for a long period of time
Alpha is concentrated in few signals or few data sources
Unsophisticated portfolio optimizer
Large differences between in and out sample or live performance
Exposures to risk factors much larger than their defined constraints
Infrequent portfolio optimization compared to forecast horizon
Crude risk controls (e.g., only limit dollar exposure to stocks/sectors)
Risk model not updated frequently enough
Large discretionary overrides
Insufficient focus on careful testing (over fitting or data snooping)
Improper use of leverage

Section Three: When does statistical arbitrage work?

Sopa Piranha is a market commentary and I have just spent ten pages explaining a strategy rather than a market. I now find myself in with the hard task of covering a complex topic of the right market environment in my four page remaining limit. Experts should call for a more detailed explanation.

I. Can statistical arbitrage make money with volatility in stocks so low?

Many people assume that a surplus of trading in statistical arbitrage explains most of the decline in returns over the last three years. There is **no evidence that the strategy has more assets** as a percent of all hedge funds. In fact there have been large redemptions recently. Overall there is more statistical trading in stocks as a percent of market capitalization, but there is also no evidence that this explains the decline in returns. Other explanations include the drop of volatility in stock indices. The Dow volatility is at a 100 year low when measured as the annual high minus the low then taken as a percentage of the low. Volatility is important for statistical arbitrage but it is **cluster volatility and individual equity volatility that matters**, not the volatility of an index. These measures can be quite different.

Below is a chart showing the difference between the VIX and the dispersion of individual stocks. Dispersion is remaining above index volatility in the same manner that it did from 1993 to 1997. The chart

shows the **first signs of life in dispersion** since it began its collapse in mid 2002. This is right on cue in terms of the level of interest rates and the attendant decline in liquidity. This is an important signal given that **declining dispersion and declining volatility are bad for statistical arbitrage**. Declining dispersion is worse than low volatility. One reason is that declining dispersion occurs when the market is full of weak players. As these unwind, the players are selling into the positions of the strong players. Like any market that unwinds, the initial cost is offset by the rebound where opportunities are a bit easier to find. That rebound is occurring now.

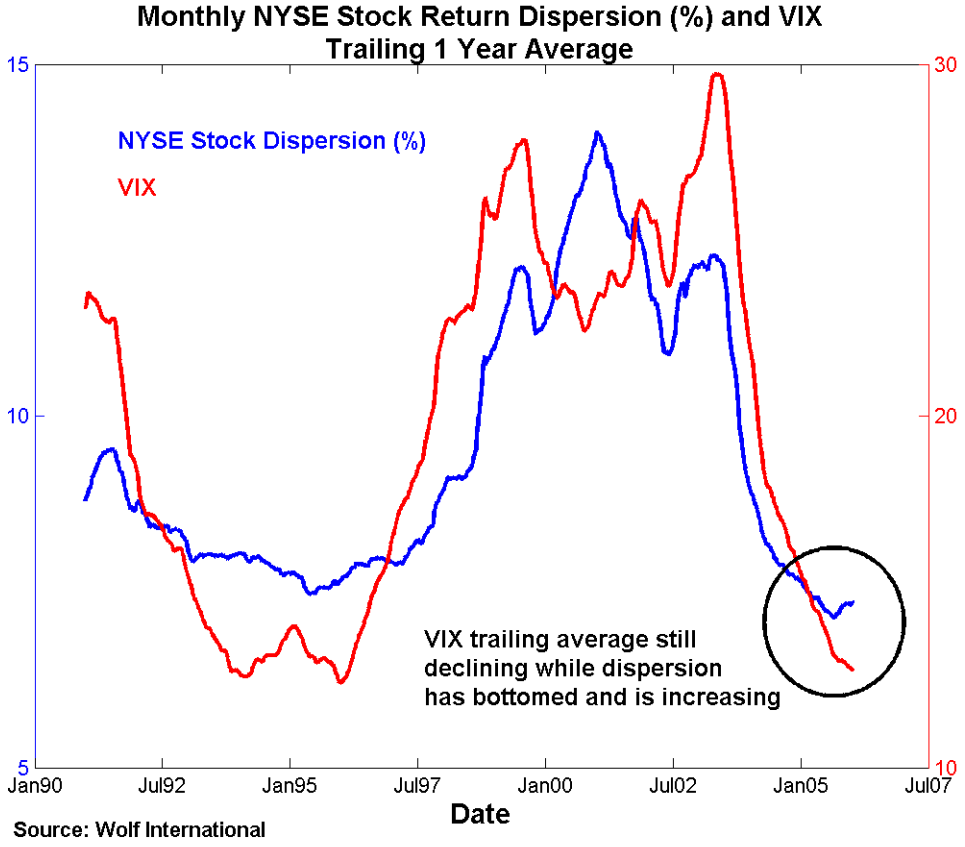

I defined dispersion as the cap weighted difference of the (daily) return of an individual stock to the market return. This is related to volatility but not exactly. Volatility can go up, but if correlations between stocks increase as well, dispersion could stay the same or even drop. This is highly unlikely, but gives a good sense of the differences between volatility and dispersion. When volatility goes up quickly, **correlation between stocks may also increase** and as a result dispersion may actually decrease. In this scenario, there is no reason for market neutral/stat arb strategies to become more profitable. High cross-correlation also reduces the effectiveness of asset diversification. What a stat arb manager wants is

volatility up and contagion/coherence/correlation to stay the same or go down.

The best environment for stat arb is high stable volatility and dispersion, second best is rising volatility and dispersion, third is low stable volatility and dispersion. This generalization holds for most stat arbs but they do vary. Some systems love rising dispersion while others struggle if it rises too quickly. For those concerned about the collapse of global monetary liquidity, it is worth noting that only **three hedge fund strategies like rising volatility**. These are; short term stat arb, ctas, and short sellers. Of those three, stat arb has the highest probability of making money. I am certainly not making a case for stat arb to be in its sweet spot. I am making a case that stat arb is bottoming now that the redemption cycle is over and dispersion is unlikely to fall. If a manager is doing pretty well now, you can be sure that when dispersion rises, the manager will really rock.

While volatility is not a necessary condition for stat arb to make adequate returns, it is a necessary condition for simple pair trading to make returns. This is why so many simple systems that are not really stat arb have been washed out of the market in the last three years. What remains are generally the best in class groups. Some remaining groups are doing well because they have high beta. True stat arb does not have much beta.

II. **Why will dispersion and volatility rise?** The Street has recently produced some research making the case that volatility will soon rise due to its negative correlation to lower global liquidity. The research shows a one year lag suggesting that now is the time for volatility to bottom. I agree though I think that many people will go broke just buying it and waiting for a big event. Investing in stat arb is a far better way to play it. A stat arb book produces a positive return with a **long options bias**. It behaves somewhat like an income producing put. Of course this does not mean that there is a real put if the stock market crashes 25%. Cheap far out of the money options can be added to the portfolio for those worried about impending doom. In the meantime, stat arb gives great diversification and high returns when leveraged appropriately.

The street research showing high negative correlation between declines in liquidity and volatility miss the direct connection through interest rates. Recall that options are a function of the interest rate, the value of the underlying instrument and the volatility of the underlying instrument. If the underlying instrument is a Company with costs that rise with interest rates, then the earnings of the company will be affected eventually. So far this has not happened because wage deflation has been an offset. More important, spreads are still extremely tight and the yield curve is flat. So the cost of debt for real companies is still quite low. Rates are not yet high enough to do much damage to companies. Volatility cannot systematically move to a higher level until rates affect companies. However, it does not

matter; dispersion is high enough to make excellent money. Dispersion seems to have made a bottom and the turn of rates is enough to predict it will stick.

Recent work at Wolf on stat arb has resulted in a close relationship to some shops that have given us an inside look at data. This data is both before and after fees. The data has allowed us to study in depth whether returns are solid. They are. We created a product combining these stat arbs to show just what is possible. During the “poor” period from 2003-2005, these stat arbs produced decorrelated returns. By combining them carefully and then adding leverage that is common in the market, we were able to duplicate the type of returns being produced in the big shops. The “group” return is 15.7% on a standard deviation of 10.5. The beta is 0.05. The Sharpe is 1.3. These are real results, not a back test. The skill added is in the combination of the portfolios. Several good models are made into one great model (very carefully). With no beta, the issue becomes change of environment and systematic risk. These were studied as well.

After all this work, we then focused on layering this over an equity index to see how the product would look (and to see why one famous manager did this on a huge amount of money). We found that this famous manager knows what he is doing. Stat arbs combined can be constructed in a way to layer over the S and P and create quite an attractive one security portfolio. Not to be outdone, I moved this product overlay to the MSCI World using a swap and discovered an even better product. The product is designed to dampen vol enough to turn the MSCI World into a 2x equity with the same volatility. The chart of the result is what you see below. The annual return is 28.0% with a 14.9 SD and a 1.8 Sharpe. Of course equities did well during the test period. If I include 2000-2003, it looks even better. This was a remarkable period for stat arb with vol and dispersion off the chart. The point of course is that they are negatively correlated. Stability of this correlation is all that would prevent any of us from buying this as a one security portfolio. I am comfortable with our work and will create it myself for half the fees of the famous one. It is a perfect addition to the fad for long only right now.

Conclusions: After a poor 2004, computerized systems are humming. Those with several combined systems are operating in a very stable way and producing huge results. The top ten stat arb shops in the world are all producing strong results on a large capital base now. We recently added to stat arb via a multi-strategy shop and one direct multi-system stat arb. We are evaluating all of our top ten manager picks and one more will be added. However, what is way more compelling is a customized product that combines some of the single system type shops. These shops do not leverage as much as they should. In addition, I am swayed by the MSCI overlay product. I would like to add it on a market dive of more than 15%.

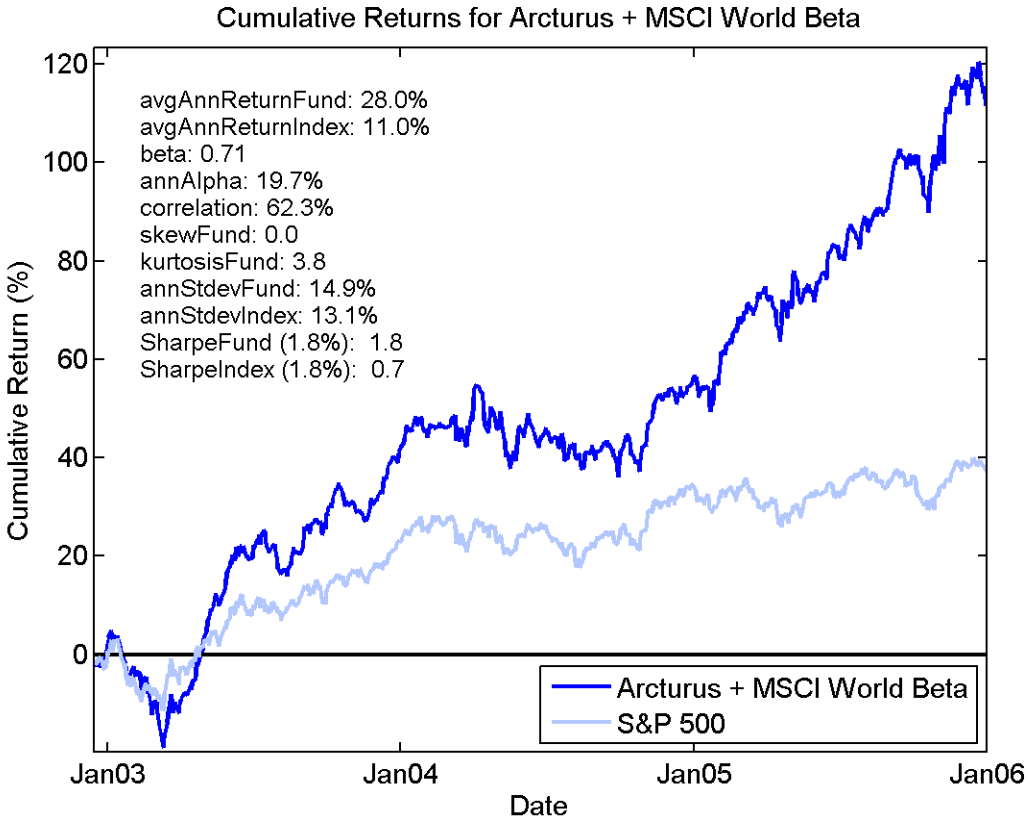

Good luck in your investing and we hope you enjoyed your piranha soup…
Mari Kooi
Wolf International

Sopa Piranha was originally published by Wolf International, 16 February 2006.