

Tall dark strategy?



Stuart Farr

What does the future hold for quant traders? Stuart Farr, President of Deltix, looks into the very long term.

Statistics keep coming. Depending on the source, we read that some 60% of US equity trade volume is a result of high frequency trading (HFT) and 25% of FX trading arises from algorithms, with both figures rising. The Aite Group estimates that 12% of global assets under management were driven by quantitative analysis in 2007, a figure expected to reach 14% by the end of this year. Whether it's investing or trading, it is abundantly clear that quantitative approaches are increasingly pervasive.

There have been two major drivers of these trends. The best execution requirements of Markets in Financial Instruments Directive (MiFID) and Regulation NMS and the consequential fragmentation of liquidity resulted in new trading opportunities for liquidity makers (but added unwelcome complexity for liquidity takers). Critically, the *enabling* driver has been the persistence of Moore's law, and its extensions, and the concurrent reduction in the cost of technology.

This has had two effects. First, it has extended the reach of quantitative trading into the budgets of increasing numbers of money management firms: institutional managers, hedge funds, CTAs and independent proprietary trading firms. Some of these firms are just two people, but they have the technological horsepower to have a profitable business with maybe less than \$20m. Secondly, it has enabled the current arms race of ever faster trade execution: the proverbial "race to zero". One can wax philosophical about the merits or otherwise of this democratisation of quantitative trading. What is axiomatic, however, is that more players equate to more competition and crowding.

It's Not All About You

There are fundamental capacity constraints inherent in liquidity-taking HFT strategies resulting from the need for deep liquidity required to minimize the market impact of quickly entering and exiting positions. With further reductions in the cost of technology, increasing competition in HFT is inevitable. This competition, and consequent reduction in rewards, is of course exacerbated as the ultimate latency, zero, is an immovable object. Liquidity-making HFT strategies face the same technology-led competitive forces but, as liquidity providers, are not constrained by liquidity themselves.

But, despite the outsized attention which HFT garners, it is not the only form of quantitative trading. What about lower frequency strategies? As we move away from HFT, we can more clearly distinguish between portfolio allocation strategies (quantitative trading) and efficiently executing trades generated from these quantitative strategies, that is, algorithmic trading. Clearly, for HFT strategies where positions are held for seconds only, such a distinction is moot. But for quantitative strategies with holding periods of days, weeks or months, then quantitative trading (as in the portfolio allocation meaning) is clearly a discipline


distinct from algorithmic execution. Algorithmic execution is, in these non-HFT strategies, a cost-minimization exercise; the objective being to lose as little alpha as possible between signal generation and trade execution.

Where Do We Go From Here?

In August 2007, at the start of the financial markets collapse, many fund managers with deteriorating subprime credit positions had to raise cash quickly to meet margin calls. The resultant sell-off of equity positions hit quant funds particularly hard. The massive drawdowns amongst several of these funds increased the speculation that many were running the same or similar positions, arising

from similar strategies. Not surprising maybe, but a disconcerting reminder of the crowded nature of some significant (in terms of capital) strategies. One of the first and still a major quantitative trading strategy is statistical arbitrage ("stat arb"). As a result of reduction of pricing inefficiencies, stat arb trading, in US markets particularly, is less profitable now than ten years and more ago.

Of course, the net result is that some quantitative strategies have lower longevity – resulting in lower lifetime profit. Clearly, arbitrage strategies operating in efficient markets are particularly susceptible to this; liquidity-making strategies are not, but are themselves more susceptible to the technology arms race.





“...many fund managers with deteriorating subprime credit positions had to raise cash quickly to meet margin calls.”

So, with fewer and less pronounced pricing inefficiencies (at least in the US equities markets), ever increasing competition and crowding in both many low and high frequency quantitative trading strategies, where do we go from here?

Expanding the Universe of Opportunities

Pricing efficiency, crowding and competition are most acute in the US and to a lesser extent the European equity markets. Consequently, some firms are looking at opportunities in other areas: geographic and asset classes.

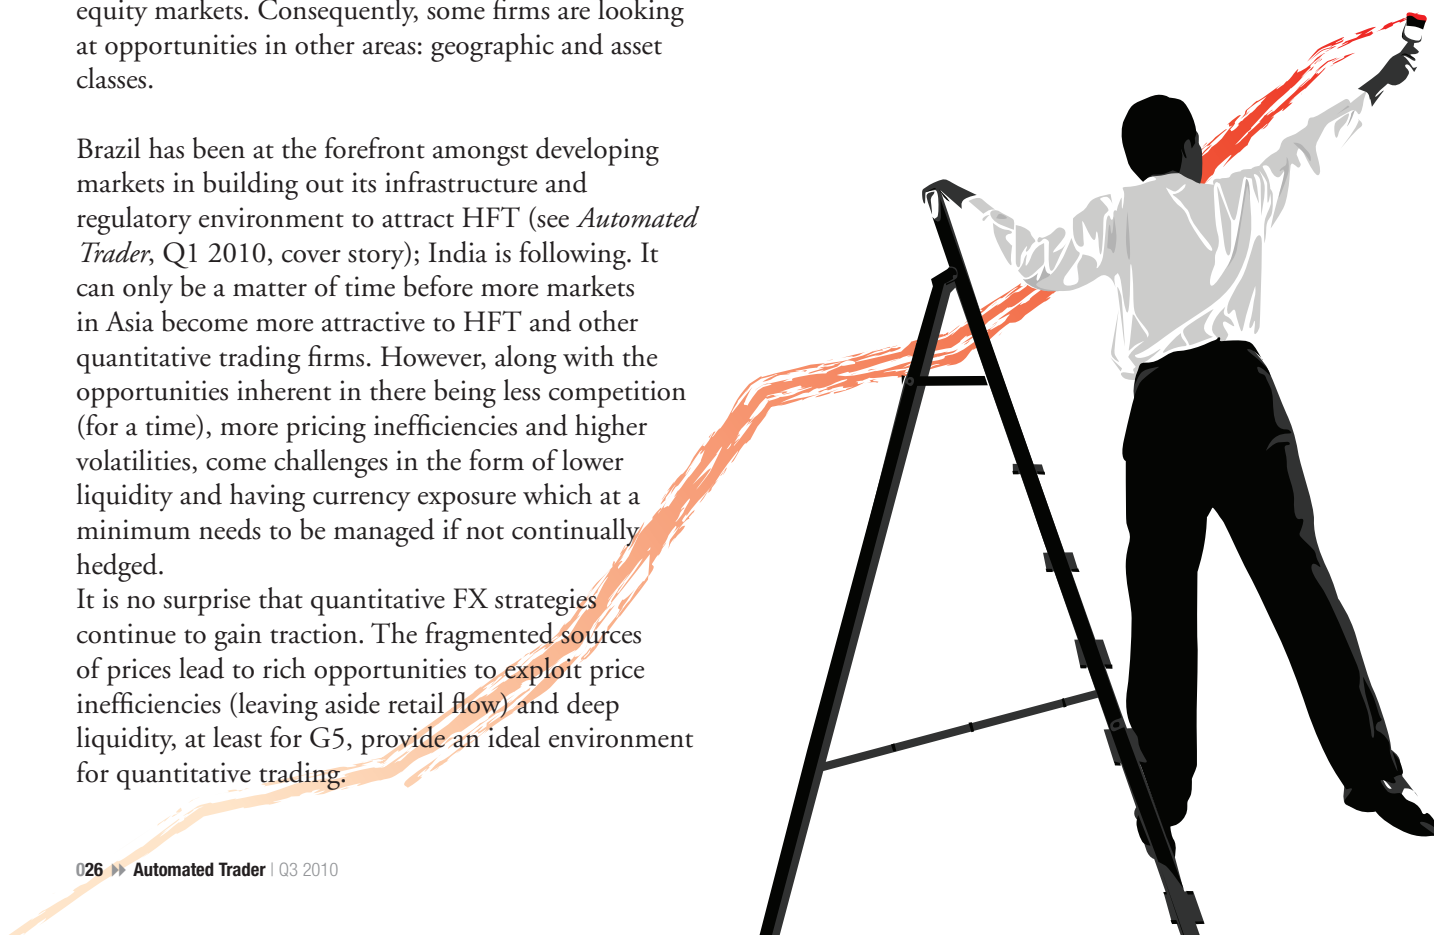
Brazil has been at the forefront amongst developing markets in building out its infrastructure and regulatory environment to attract HFT (see *Automated Trader*, Q1 2010, cover story); India is following. It can only be a matter of time before more markets in Asia become more attractive to HFT and other quantitative trading firms. However, along with the opportunities inherent in there being less competition (for a time), more pricing inefficiencies and higher volatilities, come challenges in the form of lower liquidity and having currency exposure which at a minimum needs to be managed if not continually hedged.

It is no surprise that quantitative FX strategies continue to gain traction. The fragmented sources of prices lead to rich opportunities to exploit price inefficiencies (leaving aside retail flow) and deep liquidity, at least for G5, provide an ideal environment for quantitative trading.

Multi asset class quantitative trading is, anecdotally, increasing and likely to continue to do so. There are many strategies for trading options and their underlyings quantitatively. As computer capacity and software sophistication continue to increase, the complexity and breadth of these quant strategies is increasing. The current legislative pressure to trade hitherto OTC derivative contracts on exchanges will overnight provide quant traders new data sets to analyse and new trading strategies to devise (as well as provide new opportunities for market data providers). But what about other asset classes without explicit relationships between instruments?

Quantitative trading of fixed income is currently small relative to equities, FX, futures and options. Electronic trading of fixed income is likewise less widespread than other fungible asset classes, but it looks like we are set for an increasing supply of government debt for several more years yet.

By definition, the central requirement of quantitative trading is data, particularly market data. Using fundamental data for quantitative investing and indeed



Tall dark strategy?

simple stock screening has been practised for decades. Some quantitative strategies include fundamental data as qualifiers to their price driven strategies: the rationale being to validate on a real-world basis (because financial statements *never* mislead) the proposed signal generation. In a similar vein, there is increasing use of digitized news data. Such data is also used as the sole data input for some models. The use of digitized news is still in its infancy, but its scope is very interesting to the extent that it is much less deterministic than price/volume data and it offers new areas for alpha discovery.

Doing More Of The Same: Increasing Throughput of Ideas, Increasing Quality and Reducing Time to Market

Doing more of the same does not sound as interesting as developing trading strategies for new markets or asset classes, but building from a trading firm's existing expertise is likely to be less risky, certainly simpler and inevitably incur less organisational inertia. The Aite Group has defined the quantitative analysis and trading process (also known as Alpha Generation) as a series of steps in a workflow, illustrated below.

According to Aite, these steps on average can last between 10 to 28 weeks. Given that the length of time some trading strategies can be profitably deployed is decreasing (particularly liquidity-taking strategies), then many firms have to increase the throughput of the alpha generation process. Improving the efficiency of the alpha generation process is achieved in the following ways:

1. Systematizing the process
2. Improving the efficiency of each step of the process

Systematizing the Process: Ben Van Vliet has published work, including in this magazine, on systematizing the alpha generation process. Key to this is actually defining a series of pre-requisites and steps to follow. For example, defining the people involved in each step of the process. Especially important in this regard is the decision if, and if so, how much initial capital for each strategy should be deployed. In larger firms in particular, this can clearly be a cause for contention amongst competing individuals or teams and so a defined process, complete with transparent decision making factors, for making this determination is essential.

Improving Efficiency: There are several ways of improving the efficiency of the alpha generation process. Firstly, the run-time performance of the software is the major determining factor in the length of time required to back test trading models over a given data set. Clearly, more robust back testing requires more data over more time. Further, the performance of the back testing engine (comprising the time-series data warehouse, CEP engine, model and trading simulator) will itself determine the extent to which how many strategies can be back tested in parallel. Back testing run-time performance should be measured in seconds when testing one strategy on all instruments in the S&P500 across years of intraday price data.

A key aspect of the quality of back testing is the closeness of the simulated execution of the trading signals with what have occurred in live trading. This fidelity of execution is a function of both the quality of the historical data and the trading simulator. Cleansing historical market data is an art. An art made easier by recording your own and applying known adjustments for "bad" price quotes and trades, and corporate actions. The trading simulator is a crucial



"The mere addition of single tick slippage is enough to turn a strategy with a significant positive Sharpe Ratio to an outright loser."

component in the back testing process. The mere addition of single tick slippage is enough to turn a strategy with a significant positive Sharpe Ratio to an outright loser. Fine control and tweaking of trading simulator parameters, such as slippage, transaction costs, liquidity etc, is essential.


The analysis of results after the back test should be immediate. That is, performance measures such as P&L, Sharpe, Volatility and Drawdown at a minimum should be immediately assessable. This means, in conjunction with rapid back-testing, a large number of strategies (with minimal optimization) can be assessed in a very short period of time, literally hours or days. This can be thought of as an initial screening of ideas. This initial screening weeds out ideas

whose performance metrics are (objectively and deterministically) deemed insufficient for further analysis. The resulting set of successful strategies can be further refined by optimization and further tested using different time windows, including live data, until acceptable and stable results are achieved (or not, in which case such strategies are also discarded). The remaining trading strategies are available for deployment in the production trading environment.

This cycle should of course be repeated continuously by research staff in the trading firm. With several researchers, there can be a continuous flow of approved, production ready, trading strategies.

You may find yourself optimizing on one set of market conditions while trading a completely different set of conditions

The final step of deploying the strategy into production is often, somewhat surprisingly, a cause of inefficiency and of enhanced operational risk. This occurs if the trading strategy model is developed and tested in a software language that is incompatible with the production trading system. As such, the trading strategy has to be re-written in a language and/or environment compatible with the production trading system. This clearly introduces time delays, is costly



Source: Aite Group

	Data Acquisition & Preparation	Alpha Discovery	Back-Testing	Analysis & Optimization	Simulation Live Data	Production
DETAILS	• Acquiring various types of data for model creation and testing	• Creation of initial alpha model based on analysis and manipulation of data	• Back-testing of model using various historical data sets	• Analyzing the results of back-testing and optimizing and fine-tuning the model for improved performance	• Simulation with live data and analyzing effectiveness of strategy	• Coding quant research model into production for live trading environment
AVERAGE TIME	• 1 to 4 weeks	• 2 to 4 weeks	• 2 to 6 weeks	• 2 to 4 weeks	• 2 to 4 weeks	• 1 to 6 weeks

Tall dark strategy?

and exposes the firm to potential differences in the results of the research and production versions of the trading model (either as the result of bugs introduced during this translation or differing characteristics of the different implementing technologies). The importance of deploying a strategy in production which is the same code that was optimized and tested cannot be overstated. Such a capability enables the highly desirable situation of moving selected trading strategies seamlessly from back testing, through simulation with live data, to production trading. Of course, simply increasing the production of alpha generation strategies can and will be counter-productive if such strategies are of insufficient quality. That is, if they fail to perform as advertised, on out-of-sample data sets and in live trading. This means that the amount of market data history, and the run-time performance of back-testing, should be such that many strategies should be simultaneously being tested. “Many” could mean thousands.

Pushing The Envelope: Genetic Algorithms and Dynamic Optimization

Challenging the “efficient factory” approach above, a mutually inclusive approach of maintaining strategy longevity is to adapt, or evolve, strategies in response to more recent market data. We can recognize two *types* of strategy optimization algorithms: brute-force and genetic; and two *styles* of optimization: regular and dynamic.

While regular optimization is intended to discover the best-performing strategies in a single time horizon using either brute-force or genetic algorithms, dynamic optimization (also called walk-forward optimization) uses a sliding time horizon window and repeats the discovery of best-performing strategies for each sliding window.

The back testing of strategies with dynamic optimization consists of multiple optimization cycles that are shifted in time. It could be described as a sequence of regular optimizations performed on the in-sample window with subsequent execution of

the best-performing strategies and calculating their performance statistics in the out-of-sample window. Regardless of the type of optimization algorithm, the main premise of walk-forward optimization is that the recent past is a better foundation for selecting system parameter values than the distant past. The hope is that the parameter values chosen on the optimization segment will be well suited to the market conditions that immediately follow.

A problem with this approach is that when market conditions change rapidly, we may find ourselves optimizing on one set of market conditions while trading a completely different set of conditions. In this case, the likelihood that walk-forward results will

be similar to the optimized results is diminished. To remedy this situation, the time increment and window size parameters can also be dynamically adjusted using additional market events and information.

Conclusion

Many quantitative approaches to trading are faced with increasing competition, crowding and the wider deployment of ever more powerful technology. This situation puts pressure on the alpha generation process continually to replace strategies whose performance has decayed. Quantitative traders are looking in various directions for additional

sources of trading strategies: new geographic markets, new asset classes and new sources of data. However, there is likely to be more easily attainable progress in new strategy development by improving the process and efficiency of developing models in managers’ existing markets and asset classes, including using adaptive strategy optimization techniques.

Stuart Farr is president of Deltix, Inc., which provides alpha generation solutions for professional “quants”, quantitative traders, portfolio managers and technologists. Stuart was previously co-founder and CEO of Beauchamp Financial Technology, the hedge fund software provider. Prior to that, he was head of Credit Risk Technology at Credit Suisse First Boston.

