High frequency automated FX trading

The concept of automated trading has attracted rapidly growing interest in recent years. In certain markets, such as exchange-traded futures, it has already become an everyday fact of life. In others, such as interbank spot FX, the party is just beginning to get underway.

Andy Webb

Irrespective of the specific market, a key focus for this growing interest has been high frequency autotrading. While technology advances have certainly been a factor in this, another major driver has been the inexorable decline in transaction costs.

Lower cost, higher frequency

Assume for a moment that the average profit per trade for a trading system should be at least ten times the cost of trade execution. On that basis, a fall in transaction costs from $50 to $5 cuts the minimum acceptable profit per trade by $450 - from $500 to $50. This makes it viable to deploy trading systems with a smaller profit target per trade but a higher trade frequency. Typically such systems will also be using very short timeframe prices (e.g. single ticks) as a data input and will typically be handling smaller deal sizes.

Yet sooner or later this increasing trade frequency runs up against human limitations. There comes a point when it is simply no longer physically possible for the trader to hit the keypad or click the mouse fast enough, not to mention managing the resulting positions.

This conflict has been a further factor in the growth of high frequency autotrading. Even where an automated trading environment generates fewer trades per market than a human trader can handle, it can of course replicate its actions across multiple markets and timeframes. Furthermore, it is far less restricted in the number of intermarket opportunities it can observe and act upon.

> ***…In futures markets this has prompted some technology vendors to deploy customer trading models on the broker/clearer's servers within the exchange…***

An automated system is also unaffected by the psychological swings that human traders are prey to. This is particularly relevant when trading with a mechanical model, which is typically developed on the assumption that all the trade entries flagged will actually be taken in real time trading.

This is sometimes hard for a human trader to do - and not just because they may be away from their desk when a trade signal is triggered. A mechanical trading system can experience long runs of losing trades, so a human trader contemplating placing a new order after suffering six losing trades in a row may be tempted to withhold the order. Mechanical systems often depend for their overall profitability on a relatively small number of winning trades outweighing a larger number of smaller losers, so this can be critical.

In futures markets this has prompted some technology vendors to deploy customer trading models on the broker/clearer's servers within the exchange, rather than the trader's workstation.

Furthermore, while the execution of the trading model may be automated, its design and coding are still performed by humans. Any errors undetected in the development stages will sooner or later emerge (probably with expensive consequences) in real time trading. Therefore it is essential to have a robust risk management infrastructure capable of terminating the activities of a rogue trading model that has run amok. Some automated trading environments already offer this infrastructure, with a broad range of controls that can be applied to the trading systems. The FX broker EBS has created a laboratory facility which allows customers to test their model trading algorithms in a secure environment using historical FX market data and live market rates as part of its Spot Ai trading offering.

Spreading the risk, smoothing the curve

One of the most important advantages of automated trading is ease of diversification. While diversification by market is widespread, diversification by timeframe and trading model are far less common. A high frequency automated trading environment is ideally suited to these additional diversification opportunities. In such an environment it becomes relatively straightforward to deploy the same (or differing) trading models across a portfolio of timeframes.

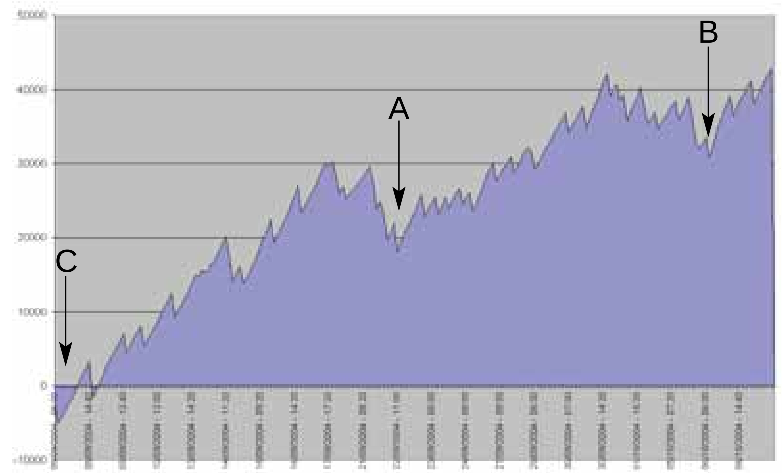

Figure 1

The benefits of time diversification in terms of reducing correlation (even when using the same basic trade system rules) can be striking. Figure 1 shows the equity curve for a very simple reversal trading system applied to 20-minute GBPUSD data over twenty trading sessions. (Twelve hour trading sessions, 0600 to 1800 UK time). Applying the same trading logic in a 1-minute timeframe obviously generates a far higher trade count than the 20-minute data (approx 5000 trades versus 270). If the 1-minute trades closest in time to the 20-minute trades are sampled to produce two sets of results of the same size, the correlation between the two data sets is just over 0.1 – a very small positive correlation. This also highlights the need for access to the most accurate and timely market data directly from its purest source.

Although the trade system in Fig 1 is profitable, it has two notable drawdowns at “A” and “B”, and an underwater period at “C” immediately after trading commences. However, by virtue of the larger number of data points it generates, high frequency trading can be added to the mix to improve the smoothness of the longer time frame's equity curve and to reduce its standard deviation of returns. Figure 2 on page 56, shows the effect of adding the 5000 results for the 1-minute time frame mentioned above to the 20-minute time frame results in Figure 1. In addition to largely removing the drawdowns at letters A and B, the

…A further advantage of automated trading when applied to very short time frames at high frequency is that it is operating in what is virgin territory…

higher frequency results also eradicate the initial underwater period at letter C. The standard deviation of returns was also reduced by a factor of seven.

Virgin territory...

A further advantage of automated trading when applied to very short time frames at high frequency is that it is operating in what is (relatively speaking) virgin territory. The physical limitations of human traders and the fact that fully automated FX trading is not universal, means that the shortest timeframes (such as individual price ticks) still contain a certain amount of non-random behaviour. At this high-speed microstructure level, the market is not yet saturated with multiple participants all chasing the same trading opportunities.

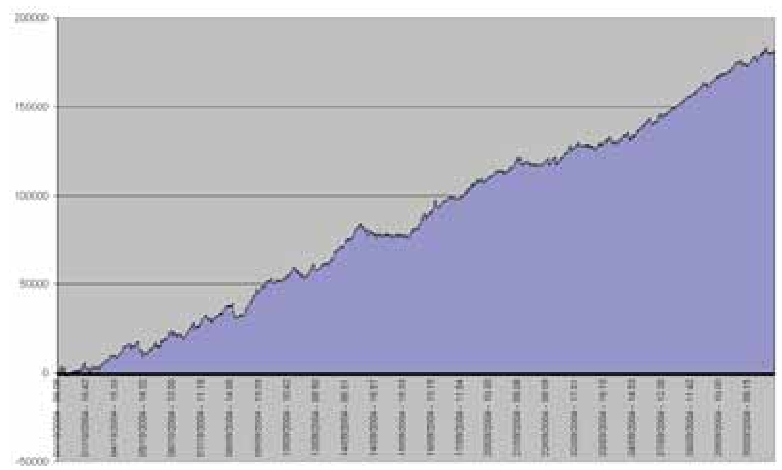

Figure 2

Figure 3 gives an indication of the sort of non-random bias that can exist at the tick-by-tick level. The "Up" column shows the percentage frequency of the market making an up tick after two consecutive down ticks and the "Down" column the percentage frequency of a down tick after two consecutive up ticks. The results are derived from testing in 4000 tick blocks at various dates, with each pattern occurring between approximately 400 and 600 times during each sample period. "Flat" ticks (consecutive ticks at the same price) were ignored for the purposes of pattern recognition.

...Commentators broadly agree that the human element will never be replaced in FX trading – there is a brain and monitor behind every mathematical model...

At first glance this degree of bias seems to offer an exceptional opportunity for automated trading. However, closer inspection reveals an important caveat that is generally applicable to high frequency trading – will an order based on these patterns actually be executable in real time? For example, a bid placed at the same price as the last down tick will be at the back of the order queue. Unless there are several subsequent flat ticks at the bid price, the order is highly unlikely to be filled. Therefore when testing such a trading model it is essential to have access to authoritative market data at a tick-by-tick level, which can be used to gauge the probability of a fill when back testing.

In addition to non-random opportunities, high frequency trading also allows existing data to be used in a slightly different way. One avenue of analysis that is often under-explored in slower frequency trading is depth of market data. For example, analysing the percentage of trades in the last fifteen seconds that have been conducted at the bid and offer and comparing that with current market depth can offer a useful indication of short-term market direction. However, this type of method is only applicable to depth of market data close to the current trading level in highly active markets. In other circumstances the data tends to be distorted by spoof orders that are "good until close".

Increasing volume...

Automated trading also has a broader benefit for the market as a whole by improving liquidity. The record volumes posted by many futures exchanges last year have been widely attributed to the wider adoption of automated trading. This is also borne out by the pilot of the EBS Spot Ai automated trading interface. All the banks participating in the pilot appreciably increased their daily volume, with the increase in some cases being very substantial.

In another case, a London based non-bank trading operation and CME customer that automatically arbitrages inefficiencies in exchange matching engines frequently accounts for 20-30,000 contracts per day in CME currency futures. During one period last year it was also responsible for more than 15% of daily volume in the CBOT Treasury Bond future.

...but increasing traffic as well

Gratifying though this surging volume is for both exchanges and participants, there is a downside – order traffic. Minimising order volume is not usually a priority for those building an automated trading system, which can result in models that fire literally hundreds of orders per second at the market. An appreciable percentage of these messages may not even be strictly necessary for the profitable operation of the model. This surfeit of orders can cause major problems for exchanges in terms of bandwidth and matching engine capacity.

...The record volumes posted by many futures exchanges last year have been widely attributed to the wider adoption of automated trading...

The need for speed - and synchronicity

This issue of excess orders is prompting a call in some quarters for controls to guard against rogue automated trading systems producing a flood of unnecessary order messages. However, especially where automated trading systems are competing for the same opportunities, the delays caused by such controls will be extremely unwelcome. In a machine versus machine trading environment, the race for competitive advantage can be determined by a few hundred milliseconds.

...In a machine versus machine trading environment, the race for competitive advantage can be determined by a few hundred milliseconds...

Although outright speed is obviously crucial, there are other equally critical but less obvious issues. One example is the manner in which events are handled within the same trading system or within multiple systems operating as a portfolio. Historically, especially when dealing with slower CPUs, there has been bias towards analytical software operating on an asynchronous basis. Under this model, the output of individual program threads is not kept synchronised, so for example one indicator in a chart window might update before another. This has the advantage of allowing efficient processor loading, and is usually of little concern to a human trader.

Figure 3	Up	Down
EURUSD	63.8	63.4
USDJPY	57.9	58.5
EURCHF	59.1	58.3

This is certainly not the case in an automated trading environment. Consider multiple trading systems applied to multiple markets, where the risk management logic requires that before a new position is opened a check is made to ensure that there is not an excessive correlation with positions already open.

For this check to be accurate, all systems must be synchronised before it is carried out. Therefore, when a global event[1] occurs (such as incoming data tick) all code pertaining to that tick in all the systems must have finished executing before the correlation check can be made. A development environment that supports this synchronous method of operation is essential if an automated trading system is to function in real time as originally intended. It also saves the developer from having to write additional code that checks that all components of a system are equally up to date before summarising and placing (or not placing) an order.

A complement, not a replacement

The opportunities offered by automated trading have prompted some to predict that it will rapidly make the human trader redundant. Those already using automated trading have quickly discovered that using it to replicate the human skill of finessing an order is difficult when it comes to particularly large or awkward orders.

As a result, the preference is to use automated trading to complement human trading activity by using it for low value tactical trading or for high frequency activity. Commentators broadly agree that the human element will never be replaced in FX trading – there is a brain and monitor behind every mathematical model. It remains the role of technology providers to give better tools to enhance their trading activity and increase their business opportunities.■

References

1 In this instance, an event that affects all the systems.

Andy Webb

Andy is a freelance writer with 20+ years' experience, most recently specialising in derivatives, technology and trading methodologies. He has written regularly for a wide range of journals - including The Sunday Times, Sunday Business, Treasury & Risk Management, Global Finance, Derivatives Strategy and Wall Street & Technology. He is also actively involved in developing and programming FX trading models for several hedge funds and proprietary trading desks.