FOREIGN EXCHANGE NEWS, ANALYSIS AND EDUCATION FOR CURRENCY TRADERS

FX TRADER MAGAZINE

BUILDING ROBUST FX TRADING SYSTEMS

CASPAR MARNEY

Building Robust FX Trading Systems

A Perpetual Quest

There has been a perpetual quest by traders, to identify quantifiable trading patterns, ever since 'candlestick' charts were developed, on the world's first trading floor; the Dojima Rice Exchange, founded in Osaka, Japan, in 1697.

The advent of computers has made that quest considerably easier, but one factor has remained constant; markets are still traded by people. As long as that remains the case, they will always be driven by only two factors, namely fear and greed. There have been many observations made about the predictability of crowd behaviour, but perhaps the most famous and often quoted is that of the famous German poet, and philosopher, Friedrich von Schiller who said,

'Anyone taken as an individual is tolerably sensible and reasonable - as a member of a crowd, he at once becomes a blockhead.'

The Holy Grail

However, although there is certainly non-random behaviour in the financial markets, equally there is almost certainly no 'Holy Grail' or 'secret formula', that even the most successful quantitative funds have discovered. If that were the case, there would be no need for them to trade so many instruments, over so many time frames, with many different models, and to focus so much of their resources on efficient execution (which will be covered in a later article).

After many failed attempts, when the author finally began to post some very consistent returns, over a two year period, a good friend inquired what his secret was and what he had discovered. He replied that he hadn't found any secret to the markets, discovered anything new, nor stumbled upon any 'Holy Grail'. He had just identified a few small, robust, edges, which were traded across as many crosses as possible, to which the very astute response came, 'That is the Holy Grail'.

Understanding the Odds

A casino only has a small edge on any given spin of the roulette wheel and will lose many times in a night, and indeed many times in a row, but statistically, it will win overall, over time:

e.g. A roulette wheel typically has 38 slots with 2 zeros. This gives the house an edge of 2/38 x 100 = 5.26%, when players bet on red or black. Even such a relatively small edge produces a substantial and incredibly consistent revenue stream for the casinos and their shareholders. The more times the wheel is spun and the more bets are made, the more the casinos probability of winning tends to 100%.

This is exactly what the systematic trader should be seeking to achieve – identifying and exploiting a small edge, as many times as possible; being the casino. Therefore, the first step in developing a robust system has to be identifying an edge. To do this, the main tools of any system developer are good historical data and software with which to analyse it. There are a number of excellent sources of data and software, readily available now. This is a huge advantage, compared to even relatively recent years, when it was very hard to come by, particularly for foreign exchange data, with no central exchange, the dominance of voice brokers and a very fragmented market. The rapid increase in computing processor speed is also a huge advantage.

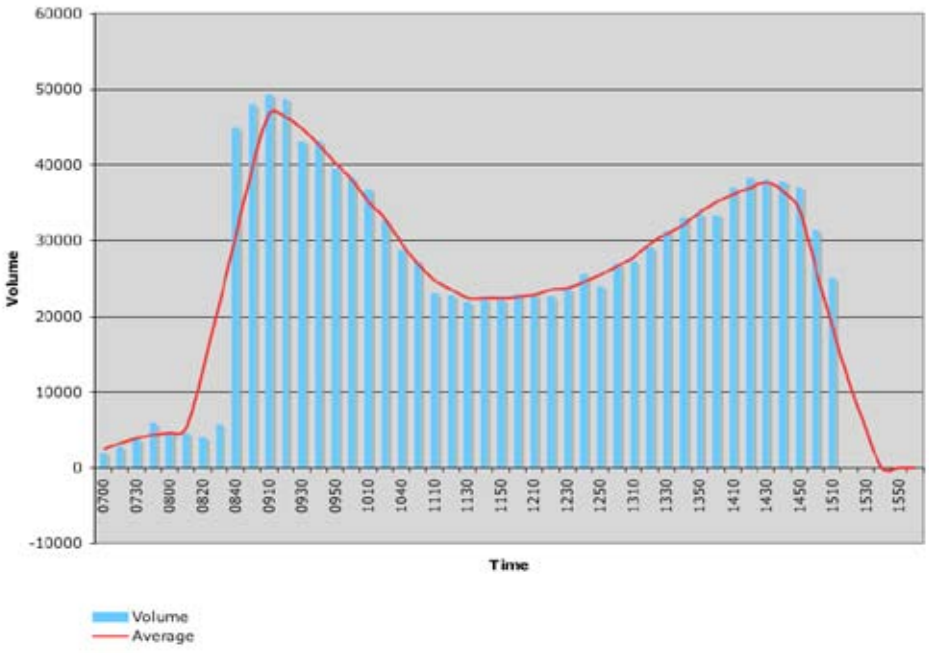

Once we have those tools in place, the next task is to quantify trading ideas and this is where any system developer will soon be able to relate to the famous Thomas Edison, inventor of the light bulb, who famously said,

'*I would construct a theory and work on its lines until I found it was untenable. Then it would be discarded at once and another theory evolved. This was the only possible way for me to work out the problem... I speak without exaggeration when I say that I have constructed 3,000 different theories in connection with the electric light, each one of them reasonable and apparently likely to be true. Yet only in two cases did my experiments prove the truth of my theory.*'

Successful Trading Systems

Unlike Edison, we have the advantage of knowing that profitable trading systems can be developed, as there are a number of proven systems already in existence, which one can easily test, such as the 'Channel Break Out' (CBO) system, made famous by the 'Turtle Experiment', where Richard Dennis and William Eckhardt had a wager about whether successful trading could be taught (and proved that it could). Those same channel break out/trend following, techniques have been exploited by many successful funds. The 'Opening Range Break

Out' (ORB) is another system, which has been proven to have a consistent edge, perhaps most famously exploited by Toby Crabel.

The reason that these systems have proven to be robust is almost certainly because there is a sound rationale behind why they work. The CBO system relies on the fact that markets trend. It has been shown that they often have larger trends than would be expected in a 'random walk' or 'normal distribution', often displaying 'fat tails'; examples of which are almost countless, with many 'Black Swan' events happening as recently as last year.

The ORB system has worked well in the futures markets, as they have a fixed open, from which to define an opening range, and all futures markets display similar volume characteristics; as illustrated by the following sample of S&P volume on the CME, taken over several months in 2008 (Local Exchange Time).
This has remained constant over time and is something that the legendary Monroe Trout also observed. In Jack Schwager's book, 'The New Market Wizards', first published in 1992, he is quoted, as saying,

'The most liquid period is the opening. Liquidity starts falling off pretty quickly after the opening. The second most liquid time of day is the close. Trading volume forms a U-shaped curve throughout the day... Generally speaking, this pattern holds in almost every market. It's actually pretty amazing.'

Foreign Exchange vs. Futures Markets

However, to develop robust FX trading systems, we have to take into account that FX behaves differently to a typical futures market, and unfortunately there is no fixed open or close; Asia is already trading as Europe comes in, followed by London. Similarly with the 'closes'; New York and Chicago are still trading, while London and Europe are going home.

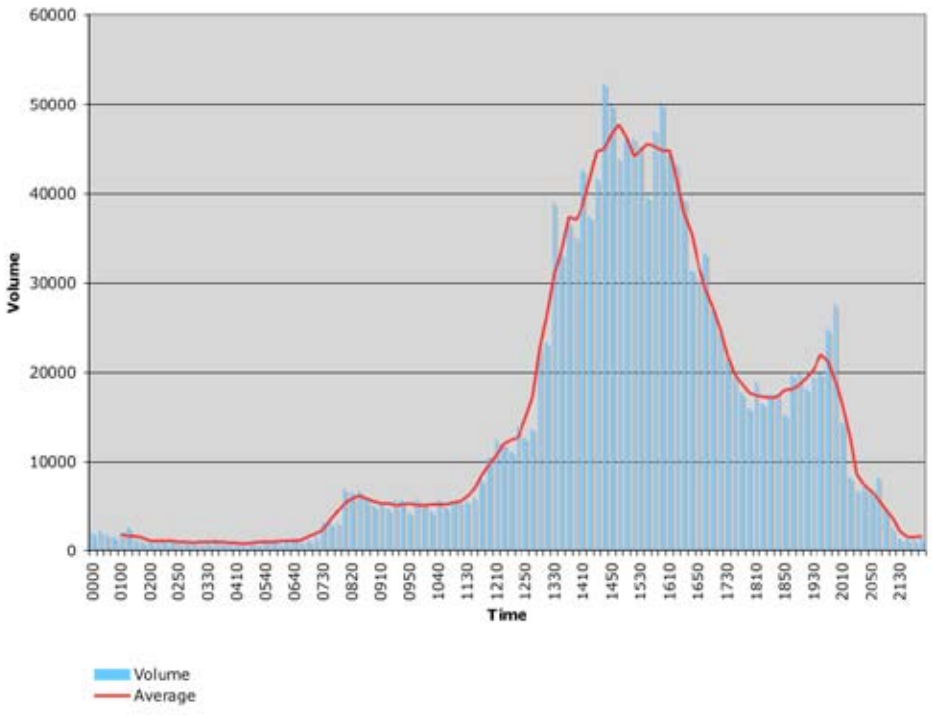

This is probably why it is considered more challenging to build successful FX trading systems: The opens and closes of futures markets are not random events and have distinct, non-random characteristics.
However, the FX markets have their own non-random behaviour. It's generally accepted among traders that each currency cross is different, with each having its own particular nuances.

To some extent, this is true, and therefore, if one finds that a certain set of parameters work well for EURUSD but not for GBPJPY, then it's easy to find arguments to explain why the two crosses may behave differently, with economic data and news events being reported in different time zones etc.

There are also moves specific to certain currency pairs, as FX is involved in every cross-border transaction across the world. As a spot trader, I recall a certain, oil oriented, corporate customer always selling a market-moving amount of GBPNOK at a specific time every Friday.
Historic price data analysis may well have revealed that non-random behaviour, but without knowing why it occurred, it would have been foolish to trade it, as one may have lost a huge amount of money if the corporate customer changed its trading habits.

There are also much broader characteristics of the FX markets. It is very well known that Europe is the largest trading centre by volume, followed by the US, with a very illiquid trading period, as the sun crosses the Pacific, until Asia comes in.

Although genuine FX volume data is impossible to quantify exactly, being so fragmented, and with no central exchange, we can use the CME currency futures as a proxy. We find that their volume distribution is very different to the distribution of a typical futures market, as discussed above. The chart above shows a similar average hourly volume for the Canadian Dollar Futures contract, over a three-month trading period (UK Time).

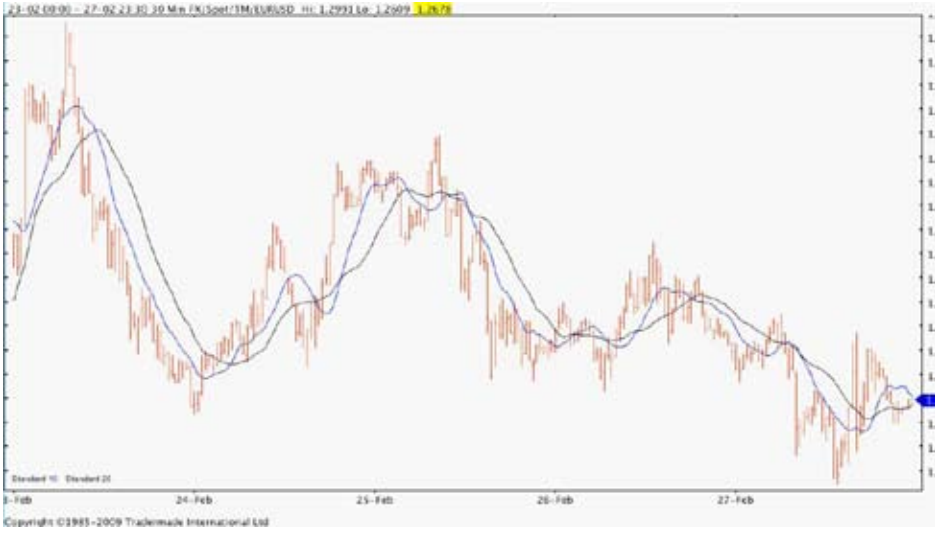

Just as with the futures markets, although volume analysis may not produce a robust trading system, it does illustrate that FX clearly isn't entirely random and there is a very predictable, robust pattern, repeated by traders every day.

Fool's Gold

With all of that in mind, it's relatively easy to find systems that work well for specific instruments, on historic data, which would appear to have huge 'edges' and to come up with explanations as to why those parameters would work for a certain cross.

When testing enough parameters though, one will always find parameters that work for any indicator on a given market. Take, for example, just testing a simple two moving average crossover combination, between 1 and 50. This will return 2,450 different equity curves (assuming we count the 10 event crossing above the 20 event moving average, as a buy, and vice versa for a counter trend trade).

By pure statistical probability, a large number of those combinations will be profitable, and statistically some of those will also be profitable 'out of sample'. In fact, it is a statistical certainty that, if you look at enough parameters, some of them will test well, both 'in' and 'out of sample'. However, without any rationale, the resulting systems would not be reliable trading systems, being solely a product of statistical probability.

Therefore, one has to be very careful and appreciate that just because a system works in simulations, it does not mean that one has discovered a robust, or even remotely reliable, trading system - another one of the countless errors the author has paid an expensive price to learn.

It is better to identify even just a small, quantifiable, edge that you understand and which has a sound rationale. To quote Monroe Trout again,

Make sure you have an edge. Know what your edge is... Basically, when you get down to it, to make money, you need to have an edge and employ good money management.'

The Edge Effect

An often-used ratio for quantifying whether a system is 'good' is 'Profit Factor' (PF), being the gross profit divided by the gross loss i.e. if the sum of all the profitable trades for a system, over a given period, was $1.1 mio, and the gross losses of all the losing trades was $1.0 mio, the Profit Factor would be 1.1/1 = 1.10

To put this in perspective, to use the roulette analogy: if a player bet on red each time, the player would win on average, 18 out of 38 spins of a wheel (with a double zero table). The house

would win 20 times out of 38 (i.e. every non-red slot). For illustrative purposes, let us assume the payout is equal to the odds. This gives the house a PF of 20/18 = 1.11[1]

With a single zero table, the house PF is just 1.05[5] (19/18).

The important point is that the house edge is a very small one, though still incredibly profitable. Looked at a different way, the odds of the house winning on any single spin of the wheel are only slightly better than evens, being 20/38x100 = 52.63% for a double zero table and 51.35% (19/37x100) for a single zero table.

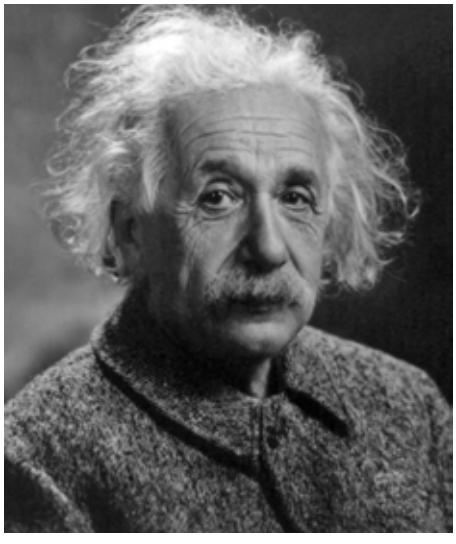

Even with only that slight edge, as Albert Einstein said,

'No one can possibly win at roulette, unless he steals money from the table, when the dealer isn't looking.'

Trading is no different. All a trader has to do, to be consistently profitable, is to identify an edge, and apply good money management. Unfortunately, that is much easier said than done. Just as the casino's edge is in knowing certain facts, a truly robust trading system, can only be built on known, quantifiable, non-random, market behaviour.

If a trading system enters a losing streak, a statistical certainty that it (often) will, it is then possible to identify whether it is just an expected statistical 'run', or whether something has changed fundamentally in market behaviour. A casino knows that each of its tables will have many losing 'runs' and it also knows that is a statistical certainty. This is where money management plays a vital role. Without understanding its edge, and without being able to quantify it, a casino would not be able to operate.

As we have seen, with an arbitrary trading system and an arbitrary set of parameters, no matter how good the 'in' and 'out of sample' results are, a system is very unlikely to be robust. Equally importantly, it would be impossible to know if the system had degraded, without understanding the underlying reason why it worked.

Caspar Marney

All of the topics touched on here will be addressed in more detail in Caspar Marney's following articles.

GLOSSARY

Base Currency: For foreign exchange trading, currencies are quoted in terms of a currency pair. The first currency in the pair is the base currency. For example, in a USD/JPY currency pair, the US dollar is the base currency. Also may be referred to as the primary currency.

Cross-rate: The exchange rate between two currencies where neither of the currencies are the US dollar.

Currency pair: The two currencies that make up a foreign exchange rate. For example, USD/YEN is a currency pair.

Forward transaction: A true forward transaction is an agreement that expects actual delivery of and full payment for the currency to occur on a future date. This term may also be used to refer to transactions that the parties expect to offset at some time in the future, but these transactions are not true forward transactions and are governed by the federal Commodity Exchange Act.

Leverage: The ability to control large dollar amount of a commodity with a comparatively small amount of capital Also known as 'gearing.'

Quote currency: The second currency in a currency pair is referred to as the quote currency. For example, in a USD/JPY currency pair, the Japanese yen is the quote currency. Also referred to as the secondary currency or the counter currency.

Security deposit: The amount of money needed to open or maintain a position. Also known as 'margin.'

Building Robust Trading Systems

Part 2 - Finding a good historical data source
"Reminiscences of a Forex Operator..."

In the last article, we introduced the idea of building robust trading systems for foreign exchange, and compared some of the characteristics of the FX and Futures markets, with FX having its own unique, but also non-random, behaviour. This article now explores the first major challenge of actually building a system, namely, building a reliable historical database:

If we were discussing futures markets, this would be relatively straightforward, as there is only one price traded at any given time with a specific volume, which is readily available, direct from almost all of the relevant futures exchanges as well as third parties. The FX market is rather unique though:

While being by far the most liquid market in the world, it's also the most fragmented. With no central exchange, each bank makes its own price, for each currency pair. Therefore, at any given moment, EURUSD may theoretically be quoted as 1.3340/42 at one bank, 1.3339/41 at another and 1.3341/43 at a third, each with their own white labelled, or proprietary, electronic trading platform, otherwise known as an ECN (Electronic Communication Network). There are also a growing number of ECNs competing for liquidity, where 'buy side' counterparties can submit their own prices into the systems. This makes it impossible to get a truly complete, clean and accurate picture of intraday FX prices. However, even the current, fragmented, electronic market is a quantum leap forwards, from only relatively recent years:

In the beginning – Voice Brokers

Before ECNs existed, most FX trading was done over the phone, with a trader sitting on a 'spot' desk, as the author once was, with half a dozen 'broker boxes', all shouting out prices. For example a 'Dollar Mark' (US Dollar v German Deutsche Mark) spot dealer (the author pre-dates the Euro) might have one broker box calling out, "thirty, thirty-five, in five" and another, "thirty, thirty-four, three by five" etc., with the 'three by five' denoting the size, in millions, that the price was good in and the 'big figure' not quoted as that was known by all involved.

Each trader, for each currency pair, would have a number of boxes shouting out similar prices and hence the classic image of a bank's trading floor, being a cacophony of sound. The reality is much different these days, with the voice brokers having been almost entirely replaced by ECNs, particularly in the major currency pairs.

In the days of voice brokers, part of the spot traders' art was to recognise the brokers voice with the best price, good in the size he wanted to execute, which as a junior dealer, was probably the hardest skill to master; particularly when the broker at each institution wouldn't always be the same person, as they would need to go to lunch, be away on holiday, or just step off the desk for a few moments. A junior on a desk would usually cover several dealers, when they similarly stepped off the desk, so may have had over twenty voices to recognise and remember which 'box' they were on. All the deals were also entered manually, unlike today's ECNs, where the deals automatically go in the trading 'blotter'. On this occasion it is probably very fair to say that junior traders today really do have it easy by comparison.

If an order was too large to execute with just one counterparty, a 'call out' would be made, where the dealer would stand up and shout, "Get me calls!" Every other dealer on the desk, would then either call up several banks on, 'The Reuters' (an inter-bank chat system) and/or the telephone. Each dealer would then shout out the prices he was being made and the dealer initiating the activity would make hand signals and shout "yours" or "mine", to indicate if he wanted to buy or sell. There was a great deal of 'spoofing' that went on, which was part of the art of good execution and mastering the art of spot trading:

For example, if a dealer at one bank took a 'call' from another, and found they were a seller, he might also sell, believing a large order was going through and expecting the price to fall, as the other bank continued to execute their order. This meant that one would often buy from the first few 'calls', hoping this would prompt the other banks to believe you were a buyer, drive the price up, quoting higher prices, into which you could then sell. Hence it was always a game of bluff, counter-bluff and spoof.

One anecdote worth recounting, in which the author was involved, is a spot desk of a first tier bank, making a huge return in the space of a few minutes, solely by a simple, but beautifully executed spoof:

The bank was known to be one that had a good relationship with the Bank of Japan (BoJ) and through which they had intervened in the market before, to strengthen their currency, occasionally coming into the market and selling a collosal, market-moving amount of USDJPY and DEMJPY. This always kept dealers wary of being the other way around, lest they got caught the wrong way on an intervention, and hence kept the Yen supported.

Therefore the chief dealer and his number two, the Yen trader, knew that if the bank was to be seen selling a huge amount of DEMJPY and USDJPY, the market may well think that the BoJ was intervening and would then also start selling, to capture the pending move down. One day they stood up and shouted "Get me calls!", which in itself wasn't unusual, as this happened on most large orders:

As each of the other traders, and assistants, all started getting prices from banks and shouting them out, they shouted, "yours!" together with the hand gesture of pushing an open hand down and away from the body (for the avoidance of any doubt as to the instruction) until they'd sold literally several hundred million US Dollars and German Deutsche Marks, against the Yen.

Nobody knew what was going on, but everyone did his or her job and got the order executed. The sales desk was asking what was happening,

as customers called up to ask what the reason was for the big move, as everybody saw and heard the huge commotion coming from the spot desk and the inevitable rumour spread that it was 'BoJ' intervention. Nobody on the desk said a word to confirm or deny the rumour, as nobody else on the desk, knew what was really going on. Just tallying the total amount sold and reconciling the now huge position the desk had, was not an easy task.

As the rumour spread and speculation mounted, USDJPY and DEMJPY continued to fall rapidly. Then came the second wave, or so everybody thought. Again the Chief Dealer shouted, "Get me calls!" and started to sell USDJPY and DEMJPY again. The market thought it was the start of a second wave of selling by the BoJ, as this was their typical style and accordingly marked their prices much lower and again sold themselves. Then came the stroke of genius – they started to buy, and buy everything, shouting, "Mine, mine, mine..." with the accompanying hand gesture of bring the palm of the hand up towards the shoulder, to the still falling prices, as other banks initially thought it was just part of a 'spoof' to sell into.

Before the market realized what was going on, they'd covered the entire position and locked in a massive profit, literally in the space of a few minutes. Everybody on the desk was given a slice of the pie, for a job very well done and it's the type of trading that we will unlikely see again – such were the days before the dominance of ECNs.

There is of course a point to this anecdote of course, other than to record it for posterity:

Although a huge amount of

transactions went through in those few minutes, none were recorded by exact time. The author himself probably executed trades, with more than half a dozen banks, but the most that would have been recorded was either a conversation on 'Reuters' or a hurried scribble on a deal ticket after a phone transaction, later reconciled with the counterparty.

Therefore, although an extreme example, it illustrates the point very well; there simply isn't a completely reliable source of accurate, historic FX data available before the dominance of ECNs and the situation hasn't improved significantly since:

The Advent of Electronic Trading Platforms

As electronic platforms began to dominate more and more of the volume, so accurate data has become more readily available, as computers are easily able to capture the exact time, price and volume of every trade. However, there is still no central ECN and rather than one becoming the dominant player, as some expected, the market has continued to fragment. This means that at every minute of the day, each currency pair is trading at different prices, bid/ask spreads and volume.

Only if one could aggregate all of the prices made on every ECN and by every bank and broker, could a truly accurate record be built. Even then though, a bank may provide a rate on several ECNs, good in $10mio, but as soon as one of its prices is hit, it will immediately 'pull' that rate from the other ECNs. Therefore, even though a 40 bid may appear to be good in $50mio, if one could aggregate all of the prices at a given moment, the reality is, that it may well not be case if you tried to execute a trade of that size.

Trading the Crosses

If someone wanted to sell the Swiss Franc against the Japanese Yen, as it's not a commonly quoted pair, it has relatively little liquidity on the electronic platforms and as a consequence has a wider price. However, USDCHF and USDJPY are more actively traded, so a professional trader would go 'through the legs' or 'components', buying USDCHF and selling USDJPY, with the USD amounts netting out to zero, leaving a CHFJPY position. This means the trader actually traded CHFJPY, but no price may actually have traded on any ECN or with any broker 'direct' in CHFJPY.

Gaps and Spikes

Although FX is by far the most liquid market, there are still times when no prices are recorded for periods of time, particularly during the less liquid Asian session and, as we have seen above, particularly in the less liquid crosses. This means that not only do genuine gaps occur in historic data, but there are also often times when a certain pair traded on one electronic platform, but not on others. These gaps in the price data need to be 'filled', which can be done using a simple algorithm, otherwise any indicator, even a simple moving average, would have an input of zero for the price at that time, which would of course create a hugely incorrect reading, which may well trigger an erroneous trading signal in an historic simulation. Conversely, not only are there times when there is no price, there are times when a spike in the data appears:

This can be due to a number of factors, but often where somebody has left an offer to take profit at, for example, 1.2580 overnight. If somebody else has a stop order to buy if 1.2520 is traded and there are no other prices in the system until the 1.2580 offer, then that would be the next price dealt. It's market practice to cancel these deals the following morning, when an obviously 'off market' rate was traded, but nonetheless, it will still often appear in the historical data made available and there is a 'grey' area where it is questionable whether the rate dealt was 'off market', or fair given the time of day and liquidity.

One of the challenges of using a simple algorithm to clean the data is that some genuine market moves can look a lot like a 'spike' in a fast market, when some news, or economic data, has just been released. A way around this is to confirm the rate via the other components. Looking at AUDUSD and USDJPY components at the time could check for example, a 'spike' in AUDJPY.

Highs and Lows

One of the most commonly asked questions in FX trading is where the highs and lows were, as this is where queries occur and money is lost and made on orders. If an order to buy was placed at 0.9840 and the low was 0.9839 offered, then the order would be filled. If the low price quoted was 0.9840/43 but was never traded, or 'given' at 0.9840, then the order would not have been filled. As it's often hard enough to determine in a real trading situation whether an order should have been filled, it's impossible to be certain with a historic simulation. In fact, if a large buy order had been placed at 0.9840, this could affect the price action itself, with market makers buying ahead of the 0.9840 bid, knowing the market will be supported there.

With the market so fragmented, and with no central exchange to determine the definitive highs, lows and the volume they traded in, order fills remain a cause of much debate, on a daily basis, in the FX market.

> there simply isn't a fully reliable source of accurate, historic FX data available before the dominance of ECNs

Predictive Pricing

As there is no central price for a currency pair, a bank or broker is free to make whatever price it wants to their customers and the customer is equally free to trade on that price, or trade elsewhere.

Some traders are very predictable in their trading behaviour and only trade with one counterparty. This leaves them open to 'predictive pricing' algorithms. For example, if some traders sold USDJPY earlier in the day, then it's likely that their next trade in USDJPY will be to cover that position and buy. Some ECNs therefore have the ability to show each customer a different price.

Therefore while a neutral price in USDJPY may be 98.94/96, one customer's ECN might show a price always marked a point higher at 98.95/97, until they have closed their short position, when it will then go back to a neutral price, earning the bank an extra pip on that trade and the customer believing he's being shown a relatively tight two point price all day.

The author has first hand experience of such pricing engines, with one of his former colleagues having built just such an engine, for a first tier investment bank. It's a perfectly legitimate practice, as the customer has the freedom to

trade on the price, or not, but the phrase, 'Caveat Emptor', is just as true in today's FX market, as it was in Roman times, when the phrase was first coined:

Consider a system, which generated a trading signal, in USDJPY, just once a day, for 252 trading days a year. Giving one point away on each trade may well result in an otherwise profitable system, recording a net loss. Without knowing why the losses were occurring, the trader may believe a perfectly robust system was no longer performing and even worse, if he were to run a simulation on that years' data, he might see that he should have made a profit, still not knowing where the 252pt 'loss' was made. This highlights how critical efficient execution is, no matter how robust the back testing and how clean and reliable the historical data; something that we'll look into in much more depth, in a future article.

Interest Rates

One extraneous factor we have to take into account, when dealing with FX, which Futures traders do not have to account for, is the interest rate differential. As each currency yields a certain rate of interest, then one earns interest in the purchased (long) currency and pays interest in the sold (short) currency. This means that if a position was held long "Kiwi Yen" (NZDJPY) then the interest rate, or 'carry' would be approximately 3pct per annum, at current rates. That is to say, if the exchange rate and interest rates remained the same in one year's time, then the trade would yield a 3pct return, being the interest rate differential earned by holding the New Zealand Dollar vs.

that paid borrowing Japanese Yen.

That difference is accounted for by 'rolling' the position every night, or 'tom/next' as it's called. When a trade is rolled, it's closed out at an agreed rate at the end of the day (called a 'reval.' being short for 'revaluation') and re-instated with a small adjustment made in the price, to account for the roll (the difference in the interest rates).

For an intra-day trading system, this isn't a factor; if the positions are flat overnight, then there is no 'roll'. For a longer term trading system, which holds trades overnight, then the interest rate differential has to be taken into consideration, to correctly calculate the results. With some historic interest rate differentials being very large, this can make a dramatic difference, and again be the difference between a system being profitable or otherwise, hence the 'carry trades' which seek to exploit exactly those differentials.

However, although the central bank rates may be fixed and known, the counterparty will usually charge a small mark-up on the 'tom/next'. Sometimes, this can be as much as several percent. Therefore it's important to know both the interest rates and the mark-up from the broker, to negotiate them as low as possible and factor them into any simulations.

Time Zones

Probably the most overlooked factor when dealing with FX data is that Europe, the US and Asia, all operate on different time zones. If we wanted to code an 'opening range break out' system for the London open, which is one of the most liquid times of day, then we should use local time in London and not GMT.

Although most data is provided in GMT, traders and therefore market behaviour, operate on local time, so daylight savings need to be taken into account. Unfortunately the US has slightly different dates when they observe DST and most of Asia doesn't observe daylight savings at all, so it's impossible to make one universal adjustment for local time across all the sessions and days of the year.

Therefore one either has to adjust the data, to local time, for the session one is interested in trading, or write an adjustment into the code, dependent on both the time of day, and date that the order is being executed. For example, if the data is in GMT, then closing a position at the close of the day in London, at 5pm, would still be 5pm in November as local time is

GMT. However, if the same trade were done in June, closing at 5pm local time in London, would be 4pm according to the time stamp of the data, as Daylight Savings would have been in effect.

Synthetic Prices

We know that the major currency pairs are the most liquid, with the better pricing; being those traded against the Dollar and the Euro. Therefore, if we had the data for those, then we could derive the 'cross rates', such as CHFJPY, GBPCAD etc.

The one challenge here of course is that if we had 60 minute OHLC (Open, High, Low, Close) data, for each hour of the day, and calculated the implied 'crosses', then we would only know the open and close accurately for those hours, as we have no way of knowing that the high or low of each component occurred (and almost certainly didn't) at the same time, within the hour.

However, it's certainly one viable method to create a reliable database. If we had hourly data for the seven major currency pairs i.e. 8 currencies, then we could calculate a synthetic price from those 'components' for the other 21 'crosses'. For example, GBPJPY is the GBPUSD rate multiplied by the USDJPY rate etc.

This creates a relatively clean set of data for the crosses, but only a line chart, as the crosses would not contain accurate highs and lows. i.e. one could not plot a bar chart, which would require the Highs and Lows of each hour.

Conclusion

Historic FX data is an absolute prerequisite, before even attempting to build a robust FX trading system, but it can only ever be an approximation, unlike the futures markets, which have a central exchange and no interest rates or 'rolls' to take into account and where the data is almost always supplied in local exchange time.

The FX market is simply too fragmented to have one universally agreed set of historic data and the trend is for the market to become more fragmented and not less so, with new electronic platforms being released each year, some carving a niche in certain currency pairs, or time zones.

Historic price data before the dominance of ECNs is much less accurate than more recent data and is, at best, an average rate traded for a certain time period. Accurate Open, High, Low and Close (OHLC) data simply cannot, and does not, exist before the days of ECNs and since then (approximately late 1990's onwards) it is far more accurate and more readily available, but can still only be an approximation. (Daily data is much more accurate as the 'OHLC' rates for a given currency pair on a certain date are generally agreed, particularly for the 'majors').

As an algorithmic FX trader, the best solution is to find a good source of data and then 'clean' it as much as possible, cross referencing the crosses and majors, filling in any gaps and cleaning out any spikes. Then the data must either be offset to account for 'daylight savings' in the time zone one is interested in trading, if the system has any time input, or it can be written into the code of a system itself.

Finally, if it's a system that holds many positions for a number of days, or frequently overnight, then the rolls must be factored into the simulations. There are many pieces of software available for analysing futures markets that can be adapted for FX data but none 'off the shelf' to date provide, as far as the author is aware, the unique functionality required to account for such unique nuances of FX.

All of these challenges probably contribute to the relative lack of successful systematic traders in FX, given its huge liquidity and clear capacity for systematic trading.

Better software and data will certainly be more readily available in the future, as FX continues to grow as an investment class. The author himself is currently involved in beta-testing a number of software packages and working with one software company to provide the unique functionality needed to test FX systems, 'off the shelf', so it's certainly something that will be available, in the near future.

In the meantime, the following resources may be useful:

Historic FX Data

Olsen Data
www.olsendata.com/

EBS
www.icap.com/markets/foreign-exchange/spot-fx.aspx

Tradestation
www.tradestation.com

Interest Rate Data

Pinnacle
www.pinnacledata.com/

Caspar Marney

Building Robust FX Trading Systems Identifying an Edge

'To succeed as a trader, it is absolutely necessary to have an edge. You can't win without an edge... incidentally, if you don't know what your edge is, you don't have one.'
Jack Schwager

In the previous two articles in the series, we discussed the need to identify a robust edge and that it must be easily explained, with a sound rationale and that it needn't be a significant edge, to produce incredibly significant and consistent returns. Just as a casino's edge is very small, when exploited many, many times, the net result is incredibly profitable.

We then discussed the need to have good, clean historic data, with which to test ideas, as inaccurate data with gaps or spikes, could easily lead to misleading or wrong results.

In this article we build on those foundations and explore the development of some ideas, from conception, through to creating trading rules, testing them and determining whether they give us a robust edge.

SUBJECTIVE ANALYSIS AND HIGH SUCCESS RATE TECHNIQUES

When I first became a trader, it never ceased to amaze me how subjective the vast majority of analysis was. The number of ideas that are in common use, many of which can be proven to be flawed, or cannot be objectively tested, upon which millions is risked daily, is nothing short of astounding.

Read almost any technical analysis on the market, easily accessible via a quick search of the web and one will find countless examples such as, 'the oscillator is overbought and therefore the market is a good sell here', or 'the market has breached the 10 day or 200 day moving average', or 'the price is at an extreme level, testing the lower Bollinger Band'.

The reason that most of these views continue to be followed is summed up beautifully by the legendary William Eckhardt, of the famous Turtle Trading Experiment,

'Since most small to moderate profits tend to vanish, the market teaches you to cash them in before they get away. Since the market spends more time in consolidations than in trends, it teaches you to buy dips and sell rallies. Since the market trades through the same prices again and again and seems, if only you wait long enough to return to prices it has visited before, it teaches you to hold on to bad trades. The market likes to lull you into false security of high success rate techniques, which often lose disastrously in the long run. The general idea is that what works most of the time is nearly the opposite of what works in the long run.'

The amount of books which also teach these 'high success rate techniques, which often lose disastrously in the long run' is equally astounding. Let us take one of literally countless possible examples from one of the better known trading strategy platforms of a Bollinger Band strategy:

Bollinger Bands LE (Strategy)

Input information

Name	Type	Default	Description
BollingerPrice	Numeric	Close	A bar price or other value used to calculate the venter-line average.
TestPriceLBand	Numeric	Close	Triggers placement of stop order at LowerBand when this price crosses over LowerBand.
Lenght	Numeric	10	Number of bars used to calculate the Bollinger band.
NumDevsUp	Numeric	2	Number of Standart Deviations for the Bollinger Band Calculation (enter a positive number; the strategy will calculate the lower band).

Usage

Long entry based on the low price crossing above the Bollinger Band.

Description

Bollinger Bands are generally placed two standart deviations above and below the market. Prices within the standart deviation are said to be ’normal' prices. Whenever the price moves below the lower band, the strategy generates a buy stop order for the next bar when the low price of the current bar has crossed back above the lower band. The stop value is the level of the lower Bollinger band.

You can change the number of bars and standart deviation used to calculate the Bollinger band.

'Whenever the price moves below the lower band, this strategy generates a buy stop order for the next bar when the low price of the current bar has crossed back above the lower band.'

This is a good example of a 'high success rate technique', which can often, 'lose disastrously in the long run'. If we applied both the long, and equivalent short, rule to AUDJPY over the last 10 years, we can see that it was indeed a 'high success rate technique', which then lost disastrously from June '08-June '09, as shown by the chart below and the equity curve in the sub graph.

However, most people trading such a technique may well believe that they were just 'unlucky', rather than appreciating the statistical certainty that it was only a matter of when, and not if, the strategy would 'lose disastrously in the long run'.

Another one of the mistakes that one sees time and time again in testing strategies is optimising the markets and parameters used. While back testing, one will find many markets where a given strategy has performed well and it's therefore a trivial exercise to construct a successful back tested simulation, of various markets and strategies that have performed well in the past.

Victor Sperandeo underlines the same point in his book, 'Trader Vic on Commodities',

'Any system or method based on optimization will fail in the long run. This is because markets change and evolve, they do not remain constant. So if you structure a system based solely on the past, it cannot survive the future.'

As highlighted in the previous articles, any trading rule will have periods and markets where it is profitable, even buying on a full moon and selling on the following full moon, will doubtless work in some markets, over some time periods. Suffice to say, that does not make it a robust strategy.

There are countless other subjective strategies which have huge followings and again are usually high success rate techniques, which therefore appear to be profitable but are possibly flawed in the long run. Many of these enjoy the benefit that they can never be disproved, lacking objective rules with which to test the theories, such as the infamous Elliot Wave or Tom DeMark studies. Though many have tried to write rules for them, I have yet to see a successful and robust translation into an objective and profitable trading strategy, though I would be delighted to do so.

ROBUST STRATEGIES

So, what do we mean by a 'robust' strategy? The foundations for a robust strategy are in having an edge and knowing what that edge is, put very well by Jack Schwager of 'Market Wizards' fame.

'To succeed as a trader, it is absolutely necessary to have an edge. You can't win without an edge, even with the world's greatest discipline and money management skills. If you don't have an edge, all that money management and discipline will do for you is to guarantee that you will gradually bleed to death. Incidentally, if you don't know what your edge is, you don't have one.'

An edge starts with a sound idea and then knowing you have an edge can only come from rigorous testing (as opposed to optimisation) of that idea, so let us start with the idea that, 'in the longer term, markets trend' and as long as markets are driven by people, fear and greed will always play a strong role and markets will therefore continue to trend.

We then need to test that idea and therefore need to develop some trading rules. This could be using a one, two or even three moving average cross over system, a channel break out, where the market makes a new 'n day' high or low, or even breaking outside of a Bollinger Band – the opposite to the strategy shown above.

The Channel Break Out system is one that has gained a great deal of press over the years, largely thanks to it being the basis for the famous Turtle Experiment by William Eckhardt and Richard Dennis. It has certainly stood the test of time and there are vast quantities of research on the system, as well as software programs, designed specifically to develop such a system, such as TradingBlox™, though it can be done in almost any software package, or even Excel.

So why has the Channel Break Out (CBO) system stood the test of time and resulted in so many successful systematic funds, when other trend following systems, such as a Moving Average crossover system, have not?

Let's analyze the results side by side. I started by taking 20 years of FX data for AUD, CAD, CHF, EUR, GBP, JPY against the US Dollar and then constructing all 21 possible crosses of those; AUDJPY, EURGBP, EURCHF etc., as described in the previous article. I also did this for intraday data, which was a considerably more demanding exercise, but am using daily data for the purposes of this analysis.

I broke the data down into two periods, 1993-2003 and 2003-2009, simply because 2003 was a convenient overlap between various data sets.

Let's start by defining the two systems:

CHANNEL BREAK OUT SYSTEM (CBO)

Buying or Selling on a new 'x' day high, or low, and closing the position out on a new 'y' day low or high. For example, if the market made a new 80 day high, we'd enter a long positions and if it then made a new 30 day low we'd exit that position, and vice versa for a short trade, as per the example below.

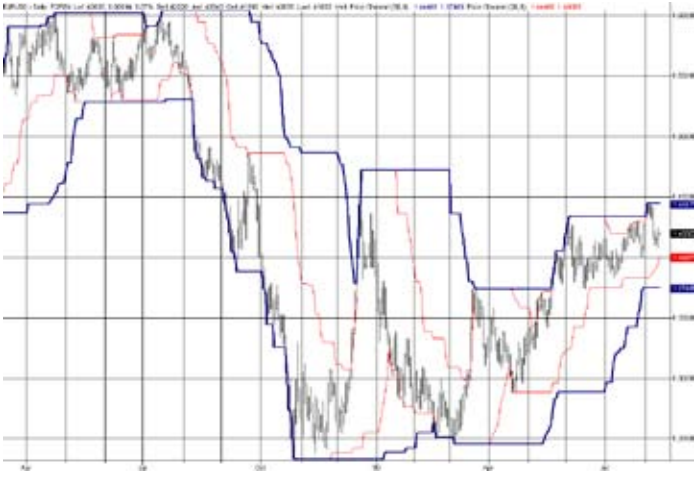

TWO MOVING AVERAGE CROSS-OVER SYSTEM (MAX)

We plot two moving averages on a chart, as per the example and buy when the shorter (fast) moving average crosses above the longer (slower) moving average:

Of course we could also trade the inverse of those two systems, selling, instead of buying on a new high, or selling when the shorter moving average crossed above the longer moving average, treating them as counter trending systems, so those tests were run as well.

We ran them in Tradestation 2000i, as that's a product many will be familiar with and into which one can easily import ASCII data files, but we could have run it in many other software packages such as Excel, Mathcad or Mathematica etc.

An exhaustive test of every CBO system and MAX system was run on each of the 21 currency pairs, over the 20 years of data, for every combination of values between 5 and 200, in increments of 5 i.e. 40x40 = 1,600 tests.

We have approximately 20yrs x 252 trading days x 21 currency pairs of daily data = 105,840 days of data. Multiply each day by the 1,600 tests, gives us more than 169 million potential trades, which is statistically a fairly significant sample.

Incidentally, this is another major mistake often made, which one see in forums all of the time; people having claimed to have found the holy grail because they found a system which performed well over the last three months on a certain instrument. This is clearly of no statistical significance and therefore such a small sample will often be extremely misleading.

ANALYSING THE RESULTS

After taking all of the results for each currency pair and converting them into US Dollars (as USDJPY produces results in JPY, USDCHF produces results in CHF etc.) we can create a 3D chart to analyze the results (using Rina Financial's '3D Smart View').

The results of the two tests are below:

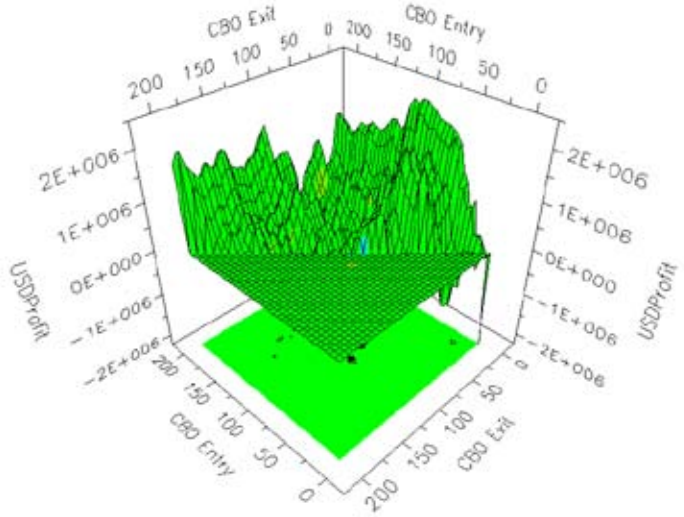

Channel Break Out 1993-2003

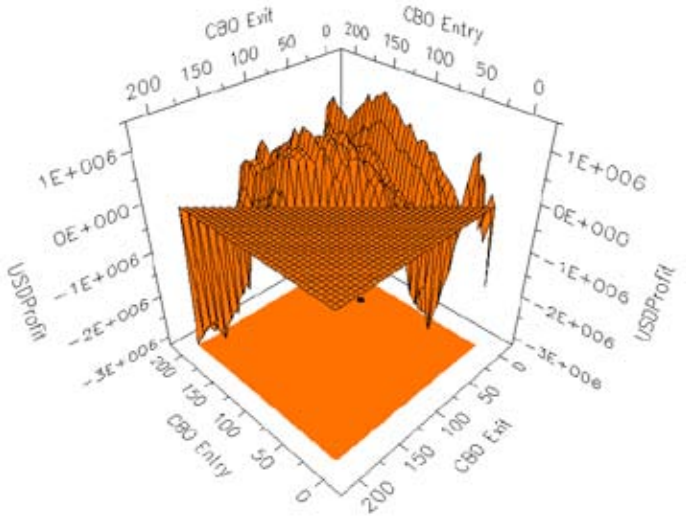

Channel Break Out 2003-2009

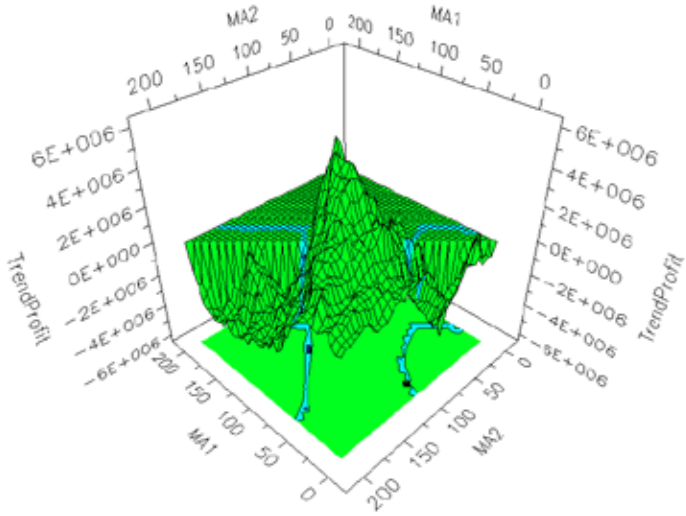

Moving Average Cross Over 1993-2003

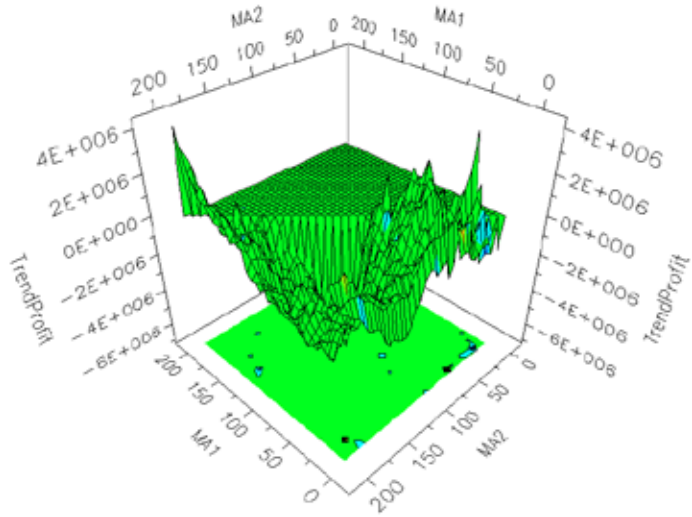

Moving Average Crossover 2003-2009

For ease of viewing only the trending half of the results are shown and what is striking is that most CBO parameters are profitable, whereas the MAX system has a distinct peak, surrounded by many losing parameters.

If the results were always stable, around that same peak, then perhaps we'd have a robust MAX system too, so now let's look at how the two systems performed from 2003-2009:

Again we see the CBO system having the majority of parameters being profitable but this time the profitable parameters for the MAX system have completely shifted to the right and the best parameters, which looked robust for the test from 1993-2003 became losing parameters in the following years.

If we look closer at the CBO system, we also see that the greater the 'CBO Entry' value, the more profitable the results. Going back to our initial premise, that for any system to be truly robust, it must be easily explained and have a sound rationale, this intuitively makes sense. The fact that a market has made a new 100 day high, is much more significant than if it's made a new 10 day high and this is born out by the result.

Also we can see that in both CBO tests that a shorter Exit signal is more robust and profitable in almost all cases, with a distinct high in the 0 to 30 day region. Again this is intuitively correct, as it allows profits to run, but cuts losses:

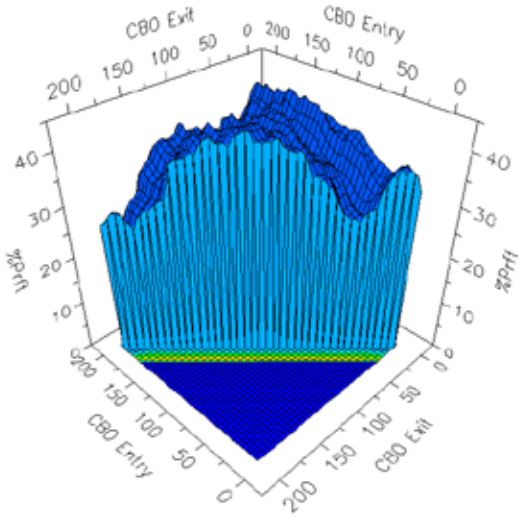

If the market made a new 100 Day high and we entered a long position, with an exit at a new 15 day low, it's going to exit the trade relatively quickly if it went against us, but it will have the ability to re-enter the long position, should the market then continue to rally and make a new high.

So let's now look at the CBO results a little closer. William Eckhardt in his interview in Jack Schwager's 'Market Wizards' told us:
'The general idea is that what works most of the time is nearly the opposite of what works in the long run.'

Above is a 3D plot of the percentage of trades that were profitable with the CBO system, using the 2003-2009 results for illustrative purposes.

Here we can see that the majority of trades are losing trades – in fact, at best, only 30-40% of the trades are profitable and this is again similar for the previous 10 years of data.

We can also look at the 'Profit Factor', which we touched on in the first article of the series. The Profit Factor is the Gross Profits of all winning trades divided by the Gross Losses of all the losing trades. For example, if all the winning trades made $1.1mio and all of the losing trades lost $1mio, we would have a Profit Factor (PF) of 1.1/1 = 1.1
Again using the 2003-2009 results, we can see that although the

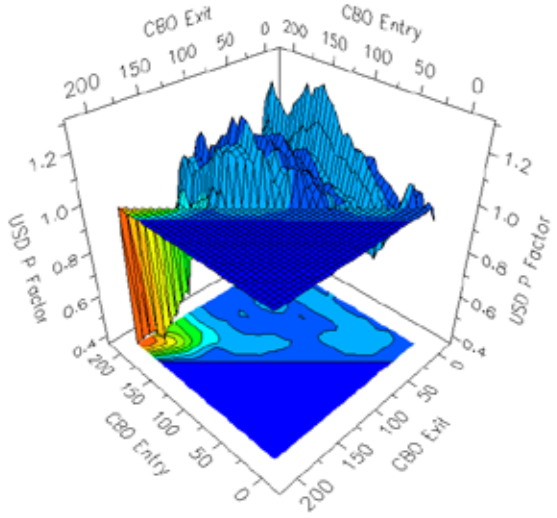

system produces robust results, the edge is 'only' in the range 1.1 to 1.2, at best. If we recall the casino comparison though, a casino's edge, when a player bets on red, for a roulette wheel with two zeros, is 20/18 = 1.1111 (where the casino wins on any black (18 slots), plus the zeros (2 slots).

By contrast, if we look at the MAX System for the two periods, we see much 'better' results in terms of both profitability and Profit Factor, with the Profit Factor exceeding 2.5 for some results in the 2003-2009 test.

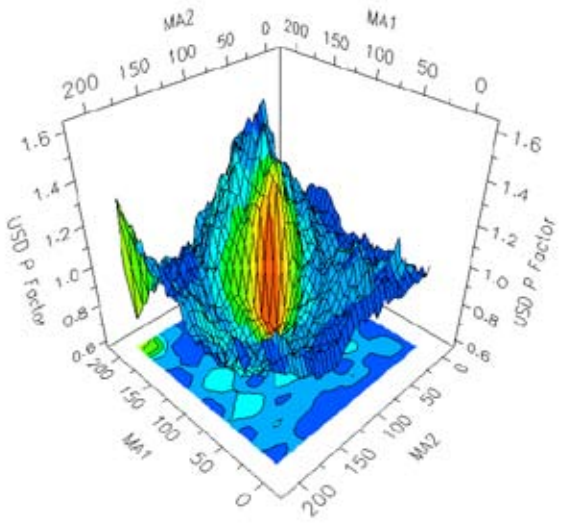

MAX System 1993-2003

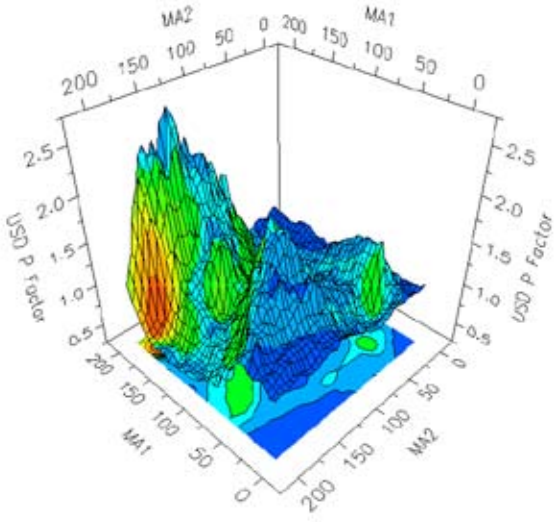

MAX System 2003-2009

However, remember that had we have chosen what looked to be the most robust results and started trading in 2003, those same parameters would have actually lost money in the following years. As Victor Sperandeo observed above,

'Any system or method based on optimization will fail in the long run. This is because markets change and evolve, they do not remain constant. So if you structure a system based solely on the past, it cannot survive the future.'

CONCLUSION

Intuitively, the results of this analysis are logical and rational, as there is very little importance, psychologically, or otherwise, that two arbitrary moving averages have crossed, no matter how good the results may look for a given currency pair, over a given time period. This is true of an infinite number of systems, as almost any system can be show to profitable over a given time period on certain markets.

This fuels the belief that systematic trading doesn't work consistently and that systems work for short periods and then stop working. That is absolutely true in the vast majority of cases, but there are clearly a number of ideas, as we have seen, which are robust, as Jim Simons (Renaissance Technologies), Monroe Trout and Toby Crabel would all certainly agree with and to which their returns stand as irrefutable testament.

When testing the CBO strategy, we have confirmed our initial theory that, 'in the longer term, markets trend'. For a robust application of that idea one would not try to pick the 'best' results from the simulations, but simply to apply some robust rules and sound money management principles.

That the market has made a new high or low and particularly a new long term high or low is important, and will likely always remain important, both psychologically and in terms of being the very definition of a trend, that the market is making higher highs.

Therefore, next time you hear someone talk about how important it is that a market has crossed a certain moving average, that the Elliott Wave is about to make an 'abc' correction, or a Tom DeMark reversal has been made, ask whether they've done the maths, and if they haven't, or have only done so with small samples, on specific markets, with limited time frames, or have optimised the results, then probably best to just smile politely, say many thanks and ask whether the market has also made a significant new high or low.

Caspar Marney

Building Robust FX Trading Systems – Exploiting the Volume Profile

In the last article in the series, we discussed robust trading ideas, comparing moving averages with a channel breakout strategy, showing how the latter is of much greater value and how using a moving average system may show great results in back testing but can be fatally flawed in actual trading.

The channel breakout strategy, while having less impressive performance statistics during 'in sample' testing, showed robust performance over time, with the same parameters providing a robust edge, over time.

The reason that channel breakout systems have stood the test of time is likely because markets trend in the long term and a new multi-month high is always going to have much more psychological significance than the crossing of two arbitrary moving averages. The findings strongly support the argument that any system based on predictable market behaviour, is likely to be much more robust than one based on arbitrary mathematical algorithms.

Therefore, in this article we are going to explore another exploitable aspect of predictable behaviour in the markets, which is much shorter term in nature; namely when traders start and end their trading day. This has been exploited in the futures markets with strategies such as the opening range breakout.

VOLUME AND TIME OF DAY

Monroe Trout, who famously made billions out of systematic trading, made some interesting observations about the futures markets when asked about the most liquid times of day, in his interview in 'The New Market Wizard' by Jack Schwager.

"The most liquid period is the opening. Liquidity starts falling off pretty quickly after the opening. The second most liquid time of day is the close. Trading volume typically forms a U-shaped curve throughout the day... Generally speaking those patterns hold in almost every market. It's actually pretty amazing."

While the foreign exchange markets have no fixed open, nor close, being fragmented between banks, brokers, electronic trading platforms and time zones, they too still display very predictable behaviour.

I have never seen similar analysis done on the foreign exchange markets before, nor seen a strategy published before to exploit the phenomenon.

This is probably because it is impossible to get accurate, historic, or real-time volume data for foreign exchange. However, it is possible to sample the market and compare the findings with other known volume information, to determine the volume profile for foreign exchange:

The EBS (Electronic Broking Service) trading platform is the largest liquidity provider and we can compare this to data also kindly provided by Barclays, from their BARX trading platform. The first two graphs below show the percentage of daily volume traded for each hour of the day, for the major currency pairs.

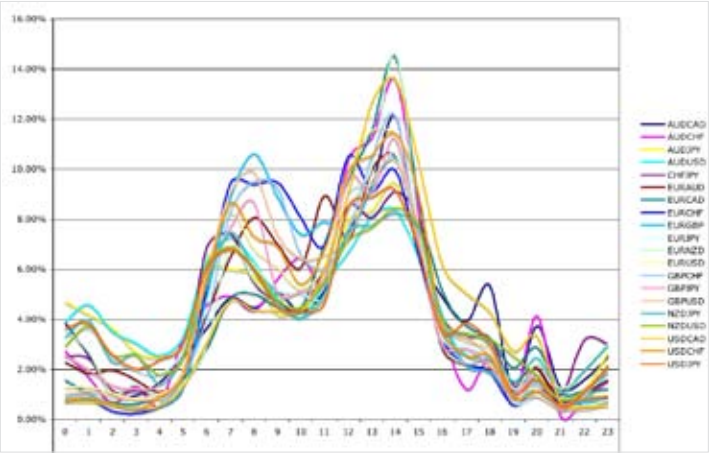

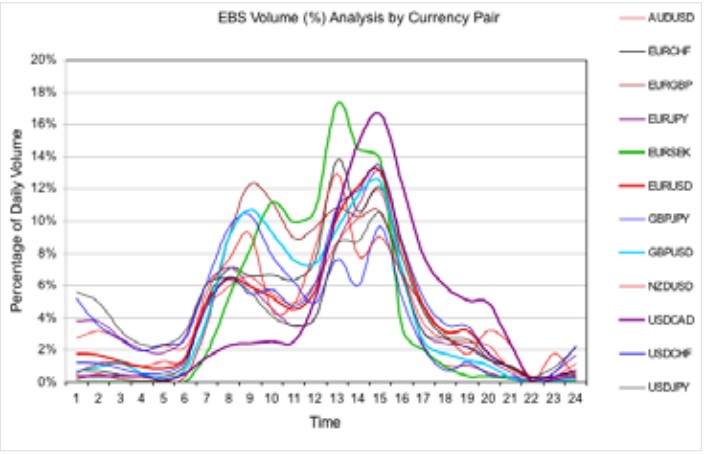

In both cases, we can see a very similar pattern emerging. Aggregating the results shows the distribution much more clearly:

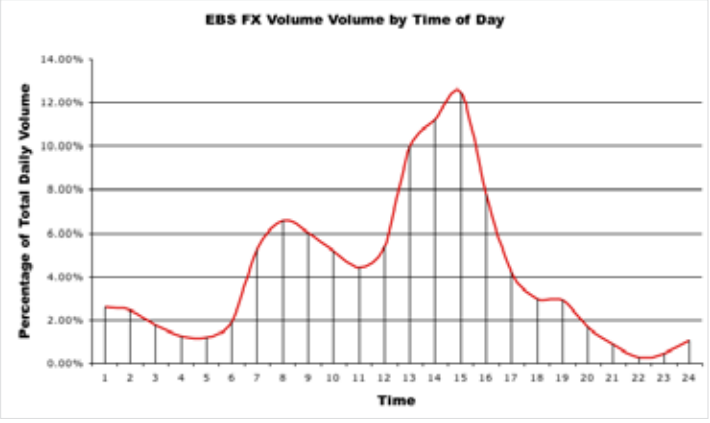

Just as Monroe Trout observed for the futures markets, although the foreign exchange markets have no fixed open, nor close and are traded twenty-four hours a day, they too follow a very predictable pattern every day.

When the aggregated volume across all currency pairs is plotted as a single graph, we can clearly see three distinct, ascending peaks of volume as first Asia and then Continental Europe, London and then the US trading sessions start. Many surveys have been done to determine the major turnover for foreign exchange by trading centre, most notably the Triennial Survey from the Bank of International Settlements, last published in 2007.

Therefore, we know that the three largest trading centres are London, Continental Europe and then the US. It's therefore not surprising to see that the largest volume of the day is during those few hours between 1pm and 4pm, during the London afternoon when the three trading centres are active. This even holds true for major Asian currencies such as the Japanese Yen that are not natively active during that time.

EXPLOITATION OF VOLUME

It is a very well proven and accepted principle of trading that volume confirms a trend. If one was locked in room, without access to any news and was only able to see price and volume, any major event would be reflected in that information. If there was a sudden move but little volume then it's unlikely that move was genuine. If however, an event such as 9/11 occurred, one would have seen both a large range in the price as well as a significant increase in trading volume. This type of volume confirmation allows a trader to know whether a move is of genuine significance.

This is one of the edges I enjoyed as a trader, while sitting on the foreign exchange desk at a major investment bank. We could physically see the customer flow going through and literally felt it, with the increase in noise. Whether consciously or not, a trader at a major investment bank cannot help but be aware of an increase in trading volume, just as a trader on the floor of an exchange is similarly aware.

It is actually almost impossible not to be aware of the interest building in a certain currency pair and this is something that almost certainly contributes to what traders often refer to as their feel, or gut instinct. It's also likely the reason why so many traders find the transition from a bank's dealing room to trading successfully outside it, to be so difficult.

One of the major challenges for an FX trader, outside a bank's dealing room, is that actual traded volume is not available in real-time across such a fragmented market, so it is very difficult to know when the volume is increasing. However, what the FX trader does know, is that any move occurring between 1pm and 4pm is very likely taking place on increasing volume, at the highest volume time of day. Therefore going with a move during those few hours is likely to provide a significant edge, over time.

TRADING STRATEGY

There are many ways to define a move, such as a change in momentum, expansion of the range, the divergence of two moving averages, the RSI crossing through the 50% level, or even standard deviations. However, let us take the simplest definition of a trend, being that of a new high or low.

We know that the lowest volume time of day is the New York close, or 10pm in London and highest volume is between 1pm and 4pm. Therefore, if we consider the New York close to be the end of one trading day and the beginning of the next, we can apply a simple trading rule to test our theory:

Buy if the market makes a new high or sell if it makes a new low, between 1pm and 4pm.

This simple strategy, without any money management, or stops, produces the following returns, based on a £100,000 account, with the portfolio being an average of the three equity curves.

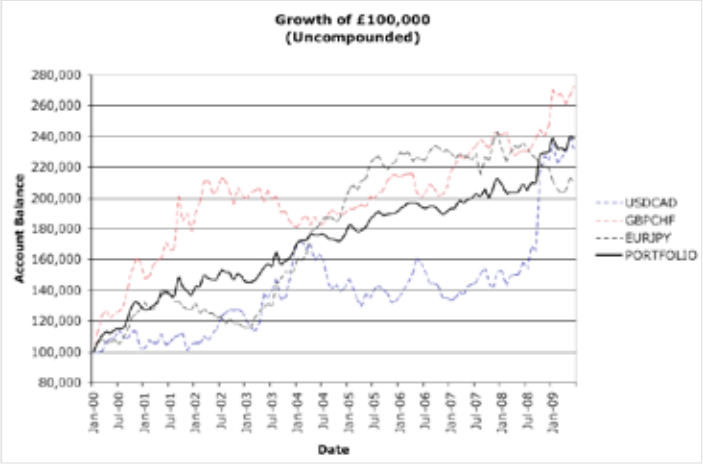

Slippage, costs and interest have not been included, as these will vary from account to account, though these factors are more than offset by the addition of some basic money and risk management principles.

We can see that the trading rule doesn't make money in all currency pairs in all years and has significant drawdowns, as well as extended periods to new equity highs. However, going with a move during the London afternoon clearly provides a robust trading edge, whether used in isolation, or as a filter to be used in combination with other trading strategies.

Again, as with the other behaviour discussed in this series of articles, we can see that while there are countless trading strategies that may work for short periods of time, based on arbitrary mathematical algorithms, there are some trading strategies that are genuinely robust, based on sound, predictable market behaviour.

Caspar Marney

Building Robust FX Trading Systems

Know Your Currencies

In the last article we looked at the commonality of volume distribution and ranges across currency pairs. We examined how similar they are throughout the 24 hour trading day, with even currencies such as the Australian Dollar and Japanese Yen having the largest volumes and hourly ranges during the London afternoon, when the three major trading centres of Europe, London and the US are all active, as opposed to their own native time zones.

Having explored the commonality, this article now explores the key differences between each currency pair, which can be broadly categorized into: time zones, liquidity, volatility and interest rates.

An understanding of these differences, that give each currency pair its unique characteristics, is important in determining whether using different parameters is 'curve fitting' or genuinely taking account of the unique and quantifiable characteristics of each market.

If the data during a back test was trending, then different parameters will likely appear to be better than a currency pair that was moving sideways during the test period. However, that is no indication that the currency pairs will continue trending, or moving sideways in real trading, so a good understanding of why a system works is vital, if it's going to be robust and continue to work in the future.

All times are expressed in local London time, unless stated otherwise.

Time Zones

We know that the London afternoon has the largest ranges and volumes for all currency pairs; however, we also know that news events and economic data affecting a certain currency will almost always occur during that currency's native time zone.

We also know that currencies are traded in pairs, so a move in GBPJPY at 2am is more likely to be a Yen move and a move at 10am is more likely to be a Sterling move. Currency pairs can therefore be broadly broken down into three categories:

1/ Currencies sharing the same native time zone:
Currency pairs such as EURSEK, EURNOK, EURCHF, USDCAD, and within an hour of each other; GBPCHF, AUDJPY and USDMXN, all share the same native trading hours.
This gives these currency pairs particularly good liquidity during their native trading sessions, particularly those whose native trading sessions occur during the London afternoon.
2/ Currencies with overlapping native time zones: Currency pairs such as GBPUSD, EURCAD, USDCHF, EURUSD, AUDUSD and NZDUSD have native time zones that overlap. Therefore a move occurring at 5am in EURCAD is more likely to be order driven, as there is little volume, news, nor economic data likely to have come out affecting that pair during those hours and would make such a move less likely to be a genuine market move.
3/ Currencies with separate time zones: Currency pairs such as EURAUD, GBPJPY and CHFJPY operate in two very distinct sessions, with their native trading hours not overlapping at all.

		Local Time in London	
Currency	*GMT*	*Start*	*End*
New Zealand (NZD)	*+13*	*18:00*	*04:00*
Australia (AUD)	*+10*	*21:00*	*07:00*
Japan (JPY)	*+9*	*22:00*	*08:00*
Hong Kong (HKH)	*+8*	*23:00*	*09:00*
Singapore (SGD)	*+8*	*23:00*	*09:00*
Turkey (TRL)	*+2*	*05:00*	*15:00*
Europe (EUR)	*+1*	*06:00*	*16:00*
Switzerland (CHF)	*+1*	*06:00*	*16:00*
Norway (NOK)	*+1*	*06:00*	*16:00*
Sweden (SEK)	*+1*	*06:00*	*16:00*
Great Britain (GBP)	*+0*	*07:00*	*17:00*
Canada (CAD)	*-5*	*12:00*	*22:00*
United States (USD)	*-5*	*12:00*	*22:00*
Mexico (MXN)	*-6*	*13:00*	*23:00*

Figure 1

If one were to express a view in Sterling at 10am against the US Dollar, or Japanese Yen, that position would be far less likely to be affected by news in the other currency, than trading against the Euro or Swiss Franc, which may have their own news, or economic data, released during that time, affecting the position. Similarly, taking a view on the Pound at 3am against the Japanese Yen is more likely to be affected by news or economic data in Japan at that time of day and has more of a Yen exposure. Therefore, any trading system must also take into account the time it is being executed for each currency pair. Figure 1 illustrates the native trading hours for each currency expressed in GMT.

Liquidity

Some pairs are more actively traded than others and this has a direct relation not only to their spreads, but also their behaviour. Although each pair may have the highest volume during the London afternoon, as we saw in the last article, some currency pairs are still far more liquid than others, throughout the day.
Figure 2 lists the currency pairs in approximate order of liquidity, with EURUSD, USDJPY and GBPUSD accounting for 52% of all FX volume (source: Bank of International Settlements, Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity, April 2007).
This makes the spreads, and therefore slippage and cost of execution, much smaller in the more liquid currency pairs. Therefore a strategy that is equally profitable in EURUSD and AUDCAD without slippage will not be nearly as profitable in AUDCAD in live trading, when spreads and slippage are factored in, due to the much wider spreads and relative lack of liquidity in AUDCAD.
It's also usually better to trade an illiquid cross via its components. For example, selling AUDUSD

and selling USDCAD results in a net short AUDCAD position. Trading via the components in this way usually captures a better spread than the trade being done 'direct' in relatively illiquid crosses.

	Turnover in Billions (USD)	*% Share*
EURUSD	*840*	*27*
USDJPY	*397*	*13*
GBPUSD	*397*	*12*
AUDUSD	*175*	*6*
USDCHF	*143*	*5*
USDCAD	*115*	*4*
USDSEK	*56*	*2*
USD/Other	*572*	*19*
EURJPY	*70*	*2*
EURGBP	*64*	*2*
EURCHF	*54*	*2*
EUR/Other	*112*	*4*
Other Pairs	*122*	*4*
Total	*3 081*	*100*
		Source, BIS Triennial Survey 2007

Figure 2

Unfortunately that also doubles the brokerage costs per million, versus being able to execute the trade 'direct' in AUDCAD, as two trades are done instead of one. This effect is even more pronounced if the strategy was executed during relatively illiquid times of day.

Further exacerbating the situation is the fact that stop loss orders can only be left 'direct', without the use of an API, usually resulting in much more slippage when orders are executed in the market.

Volatility

Volatility is a measure of how much a currency pair is moving, usually measured by taking the standard deviation of movement, over a given time period and expressed as a percentage. Traders will often refer to '1 month vol.' as a standard measure of volatility.

Essentially it expresses how much a market is moving over a given time period. A market with a low volatility is expected to have smaller moves on a given day than those with a high volatility.

For the trader, higher volatility is usually good, as profits tend to be made from movement, with a few exceptions such as some options strategies.

Volatility also has a direct impact on whether a strategy is viable. If a currency pair has an average slippage per trade of 3pts but an average return per trade of 10pts, then slippage would reduce the profits by 30%.

If the same strategy were traded on a currency pair with a twice the volatility, but the same slippage of 3pts, it may yield an average return of 20pts per trade, as the average daily movement would be higher. Slippage would then only reduce the profits by 15%; half the amount.

Slippage can be such a huge factor that some high frequency strategies, which look excellent before slippage is taken into account, can actually produce a significant loss if actually traded under real market conditions. Therefore you often see volatility filters added to strategies and these are often simply a function of the minimum volatility the market needs to be trading at, for the strategy to overcome costs and to be viable; this will vary both from currency pair to currency pair and even by time of day.

Therefore, if a strategy is found to be a losing one after slippage is added, but was profitable before, then a simple volatility filter may be all that is required, to trade only when the expected movement is above a certain amount.

Interest Rates

Interest rates are another known, and quantifiable, factor affecting currency markets. Table 3 shows the current interest rates of the major currencies.

Currency	*Interest Rate*
New Zealand (NZD)	*2.50%*
Australia (AUD)	*4.00%*
Japan (JPY)	*0.10%*
Hong Kong (HKH)	*0.50%*
Singapore (SGD)	*0.25%*
Turkey (TRL)	*6.50%*
Europe (EUR)	*1.00%*
Switzerland (CHF)	*0.25%*
Norway(NOK)	*1.75%*
Sweden(SEK)	*0.25%*
Great Britain (GBP)	*0.50%*
Canada (CAD)	*0.25%*
United States (USD)	*0.25%*
Mexico (MXN)	*4.50%*
	As of 2nd March 2010

Table 3

Trading strategy simulation software has tended to overlook the effect of interest rates on currency trading, as that is not something that affects many other markets. However, for a longer-term strategy, the effect can be particularly significant, hence the 'carry trade'.

If one were to hold a long AUDJPY position overnight, then that position would have a positive yield, or 'carry', overnight as the position was rolled. This tends to give carry trade currency pairs an underlying trend, often characterized by sharp corrections, not dissimilar to the price action of a stock market.

It's therefore vital to know if a strategy is working because of an underlying interest rate differential, as these can change dramatically over time and even invert.

Trading System Robustness

When testing a trading system, one sign of robustness is that it works across a broad range of instruments. However, when testing a currency strategy across a broad number of currency pairs, it's important to appreciate why it may show very different results and to really understand the similarities, as well as the differences between each currency pair.

For example, currency pairs with strong interest rate differentials are more likely to show trending characteristics. However, other currency pairs may trend even with little interest rate differential; the underlying reason for those trends will likely be different and needs to be considered.
Therefore you have to look at the results of any simulation to determine whether there is a valid reason for the results being different, before being able to truly decide whether a system is robust or not.
One also has to look at the price action itself and compare the equity curve of the system to the price action. It may be that the system worked particularly well in a trending market, or a sideways market, if it was mean reverting in nature.
Just because a system works on one currency pair and not another, does not mean it isn't robust. It may just mean that one currency pair exhibited a strong trend during the test period and the other did not and that may have been due to a shock news event such as 9/11, an underlying interest rate differential or a steadier shift in market fundamentals.

Conclusion

The currency markets share many similar characteristics in terms of being a global market, with similar ebbs and flows in volume and ranges, as each centre opens and close.
However, they all have their own individual characteristics in terms of time zones when news may affect that currency, liquidity, volatility and interest rates. Even though two currency pairs may share many of these characteristics, even data releases for one currency will not always occur on the same date and time as another. Therefore, at any given time, hardly any currency pairs are identical in nature but their differences are usually quantifiable.
Therefore, when testing trading strategies, all these factors should be taken into account to be able to determine whether a system, and any given set of parameters, is truly robust.

Caspar Marney

Marney Indicators

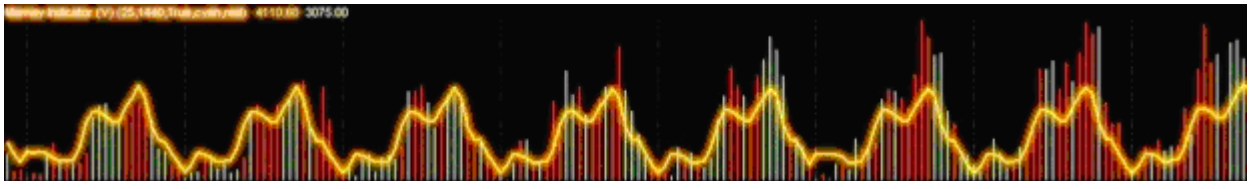

As this is the last article in the series, I'd like to introduce the proprietary Marney Indicators™, which have helped me to create profitable trading strategies by identifying and exploiting non-random behaviour.

Having learnt a great deal from other traders, sharing their insights on the markets, I hope that they will serve as a worthy contribution.

The indicators illustrate a lot of the research that has been discussed in the previous articles; the commonality of currencies as well as their unique differences, how increasing volume and range confirm a trend, as well as the importance of time as an indicator.

I was surprised not to have found these indicators already written elsewhere, as research and back testing has shown that they provide a significant edge, in exploiting non-random and therefore predictable behaviour.

Research with EBS data has shown that the number of price updates per unit of time, correlates very highly to actual volume traded. Therefore, while actual volume is not readily available for FX, many data providers now include the number of price updates, so that they can be plotted, as a proxy for volume and this can be exploited using the Marney Volume Indicator™ (MVI).

Marney Volume Indicator

The MVI plots a time-adjusted profile of volume, throughout the twenty-four hour trading day. An example of the Marney Volume Indicator™ is shown below, applied to a 60min chart of EURGBP.

I have used MultiCharts to illustrate and code the examples and Olsen Financial as the data source, as they have one of the longest historical databases available for foreign exchange, together with the number of price updates, as a proxy for volume.

The histograms show the hourly volume, via the proxy of price updates. If the corresponding hour was an up event then the bars are coloured blue and red for a down event.

The vertical dashed yellow line is a session break, showing 0000hrs GMT.

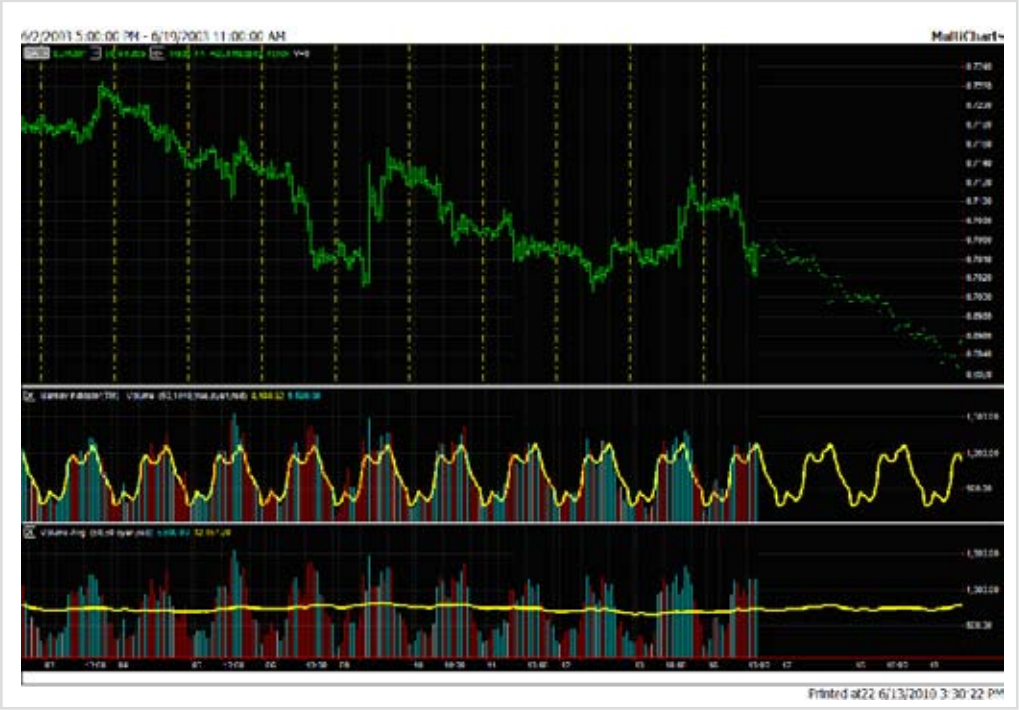

The top yellow line is the Marney Volume Indicator™ (MVI) and the bottom yellow line is a simple moving average applied to the volume, showing the dramatic difference that time-adjusting the average makes.

The MVI line shows the time-adjusted average over the previous 50 sessions. By time-adjusting the averages, the unique, predictable profiles of each currency pair are revealed. i.e. the volume from 0000hrs to 0100hrs is taken for the previous 50 sessions and the average is plotted, followed by 0100hrs to 0200hrs for each of the twenty-four hours in the trading day, by using arrays in the code for the indicator.

The Marney Volume Indicator™ therefore provides a significant improvement over the classic volume rule of simply looking for above average volume.

For any given time of day, we therefore know not only whether volume is above or below average but by how much, for that time of day also whether it is likely going to increase or decrease.

Marney Range Indictor

A similar technique can be applied to ranges, taking the true range for each hour of the day over a preceding number of days and plotting that as a time-adjusted average. An illustration is shown below using the same

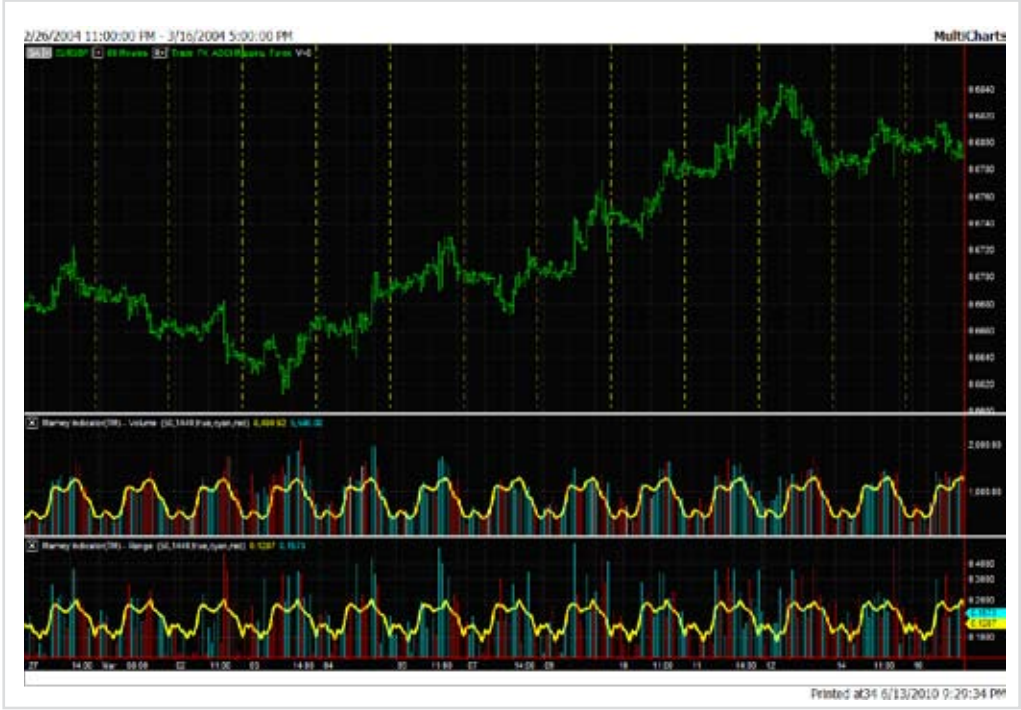

EURGBP data, showing the Marney Volume Indicator™ as the top study and the Marney Range Indicator™ as the bottom study.
As described in 'FX Trader Magazine' Jan-Mar 2010 edition, by studying historic data, we know that both hourly volumes and ranges throughout the trading day are both highly correlated and predictable. By plotting both the MVI and MRI together we can see this in real-time.

The chart below shows the indicators applied to AUDJPY.

As we might expect, the ranges and volumes are much higher during the Asian session than for a currency such as EURGBP and the peaks are much more defined when Asia, Europe and then the US enter the market.

We also see the highest volumes and ranges for AUDJPY during the London afternoon session, although not a natively active time zone for the currency pair, a common characteristic of currency pairs, previously identified and explained in this series of articles.

Real-time Volume and Range Analysis

As illustrated in the last article, profitable trading strategies can be developed from being able to predict when the highest volumes and ranges during the day are likely to occur in an individual market.

By being able to plot expected volume and ranges in real-time, those concepts can be enhanced even further. Using the MVI and MRI, we can see whether the current range and volume is higher or lower than expected for a given time of day.

If the market is making a new high and both the range and volume is higher than expected for that time of day, then the move may be considered to be more significant and conversely if a move occurred on particularly low volume and range, then it might be considered less significant.

I have carried out a considerable amount of research around this basic idea and found a number of ways that these indicators can be used, to profitably exploit predictable behaviour in the markets.

The code for both indicators is available for free from my website and I hope that it provides readers with an additional edge in their trading, whether systematically or as an additional tool for discretionary decisions.

MARNEY VOLUME INDICATOR™

```
input: avgLen(10), mins.
in.session(1440), autobars(True),
upcolor(cyan), dncolor(red);
var: start(0), end1(0), end2(0),
x(0), p(-1), count(0), avg(0),
barsinday(0), DayNumber(0);
array: xv[199,1440](0);

if bartype < 2 then begin

start= (Sessionstarttime(1,1));
end1= (sessionendtime(1,1));
end2= (sessionendtime(1,2));

value1 = timetominutes(start);
value2 = timetominutes(end2);
if start > end2 then
value3 = 1440+(value2-value1);

if start < end2 then
          value3 = -(value1-
value2);
if autobars = false then value3 =
mins.in.session;

barsinday = ceiling(value3/
barinterval);

if d<>d[1] then begin
    if count=barsInDay then begin
        p=iff(p<avgLen-1,p+1,0);
        for x=1 to barsInDay
begin
    xv[p,x]=ticks[barsInDay+1-x];
        end;
    end;
    count=1;
end else count=count+1;

if xv[avgLen-1,count]>0 then
begin
    avg=0;
    for x=0 to avgLen-1 begin
        avg=avg+xv[x,count];
    end;
    avg=avg/avgLen;
   plot2(ticks,"ticks",default
,1);
   plot1(avg,"avg",yellow,defau
lt,1);
end;

if close > open then setplotcolor
(2,upcolor);
if close < open then setplotcolor
(2,dncolor);

end;

if bartype > 1 then begin
          avg =
averagefc(v,avglen);
          plot2(v,"ticks",default
,1);
          plot1(avg,"avg",yellow,
default,1);

if close > open then setplotcolor
(2,upcolor);
if close < open then setplotcolor
(2,dncolor);

end;
```


MARNEY RANGE INDICATOR™

```
input: avgLen(10), mins.
in.session(1440), autobars(True),
upcolor(cyan), dncolor(red);
var: start(0), end1(0), end2(0),
x(0), p(-1), count(0), avg(0),
barsinday(0);
array: xr[50,1440](0);

if bartype < 2 then begin

start= (Sessionstarttime(1,1));
end1= (sessionendtime(1,1));
end2= (sessionendtime(1,2));

value1 = timetominutes(start);
value2 = timetominutes(end2);
if start > end2 then
value3 = 1440+(value2-value1);

if start < end2 then
          value3 = -(value1-
value2);
if autobars = false then value3 =
mins.in.session;

barsinday = ceiling(value3/
barinterval);

if d<>d[1] then begin
    if count=barsInDay then begin
        p=iff(p<avgLen-1,p+1,0);
        for x=1 to barsInDay begin
            xr[p,x]=truerange[bar
sInDay+1-x];
        end;
    end;
    count=1;
end else count=count+1;

if xr[avgLen-1,count]>0 then begin
    avg=0;
    for x=0 to avgLen-1 begin
        avg=avg+xr[x,count];
    end;
    avg=avg/avgLen;
   plot2(truerange,"range",defau
lt,1);
   plot1(avg,"avg",yellow,
default,1);
end;

if close > open then setplotcolor
(2,upcolor);
if close < open then setplotcolor
(2,dncolor);

end;

if bartype > 1 then begin
          avg = averagefc(trueran
ge,avglen);
          plot2(truerange,"range"
,default,1);
          plot1(avg,"avg",yellow,
default,1);

if close > open then setplotcolor
(2,upcolor);
if close < open then setplotcolor
(2,dncolor);

end;
```


Conclusion

This series of articles has been the result of years of research, learning many expensive mistakes along the way, such as identifying arbitrary mathematical algorithms that appeared to be the Holy Grail, finding systems that worked particularly well on some markets but not others and systems that appeared to work well both in and 'out of sample'.

Almost all of these ideas and discoveries were flawed.

I have learnt that each mistake was, in some way, a result of either over optimisation, or curve-fitting, even if inadvertently. I hope that these articles help others to avoid many of the pitfalls of building trading systems that it has taken me years to learn; with no doubt many lessons still to be learnt.

To summarise, in a few simple rules:

Keep it simple – if a system looks too good to be true, it probably is.
There is no 'Holy Grail' – only applying a small robust edge with consistency and discipline, over a portfolio of instruments, with good risk management.

Avoid arbitrary formula – if you test enough parameters, you will always find some that work, both in and out of sample, or on some markets. That doesn't mean they're robust parameters, nor even robust ideas.

Do base systems on market behaviour that can be explained and understood.

Remember that nothing in the world can take the place of persistence.

Caspar Marney