The Retail FX Trader: Rising Above Random

Christopher J. Davison

Nottingham Trent University, UK

February 4th 2016

Abstract

There has been much historic discussion about the effectiveness, or otherwise, of technical trading strategies in the financial markets. The view, that 'technical analysis' (TA) may be of limited value in a market seemingly driven by economic fundamentals, would seem to be supported by research showing that retail Foreign Exchange (FX) traders as a whole are not achieving returns above that of random trading. Despite this, technical trading strategies and the search for the 'holy grail' system remains ever popular with the Retail FX Trader. This paper examines three popular technical trading strategies and 'best practices' used by the retail FX trader, to try and identify rules and approaches that might help such a trader achieve 'better-than-random' trading results. Using a non-optimised, computer based trading simulation, results from over 175 million 'random' trades across nine years of data were evaluated to try and establish if such rules exist and to answer the question 'can a Retail FX Trader ever expect to use technical analysis to achieve net profitable outcomes?' The results show that the use of Technical Analysis does seem to offer better-than-random results and that setting significantly larger profit targets for trades, versus the maximum loss a trader is prepared to accept, can produce profitable trading, even when using no TA and entering trades randomly.

1. The Effectiveness of Technical Analysis in FX Trading?

There are two approaches to analysing financial data in order to make trading decisions; fundamental and technical. Whilst fundamental analysis of Foreign Exchange looks at economic and geo-political factors that might affect a countries currency value, technical analysis is concerned with the study of historic price movements, trying to identify patterns of previous activity that can be used to predict future ones, or "a set of techniques for deriving forecasts of financial prices by analysing the history of the particular price series." (Menkhoff & Taylor 2006). This might involve identifying price levels where previous buying or selling activity has taken place or identifying that a particular currency pair is 'trending' upwards. Patterns such as 'double tops' or 'head and shoulders' (Lucke 2000) are examples of technical analysis (TA). A survey carried out by the Bank of England in 1988 found that "at least 90 percent of respondents placed some weight on this form of non-fundamental analysis" (Taylor & Allen 1992) whilst a 1998 study of foreign exchange traders in Honk Kong found that ">85% of respondents rely on both fundamental and technical

analysis" (Lui & Mole 1998). Modern computerised trading platforms give the retail trader easy access to any number of charting and TA options.

The availability of technical analysis tools and the execution issues in trying to trade fundamental news means that TA is the dominant approach taken by the retail FX trader. MyFXBook (MyFXBook 2015), the web portal for retail FX traders referred to by Heiner and Simon in their study of over 5000 retail traders (Heimer & Simon 2012) shows 30% more forum posts relating to TA than to fundamental analysis, with a similar metric from ForexFactory (ForexFactory 2015), a second global portal for retail traders, shows 60% more TA related posts. However, whilst apparently being a popular approach to trading, there is conflicting data as to its effectiveness. A paper looking at the profitability of technical trading rules (Qi & Yangru 2006) concluded that "While numerous studies document significant profitability of certain rules, it is unclear whether their superior performance is due to their genuine economic information or down to pure luck". The question "do technical trading rules generate profits" was further posed (Curcio et al. 1997) with the answer being "although some profits could be made by following these rules in periods of trends, this was not the case on average". A specific study looking at the profitability of the use of head and shoulder patterns (Lucke 2000) concluded that "if there is any evidence for non-zero returns at all, then it is evidence for negative returns". Evidence of 'profitability and statistical significance of technical trading rules' (Levich & Thomas 1991) by comparing the results of random trading of currency futures contracts verses the results of systematic rules, however it is not clear if these results, given the transaction costs faced by a retail trader, would offer a profitable return. Research carried out by the author shows that retail FX traders, when viewed as a whole, do not demonstrate a better-than-random profitability (Davison 2016).

Much of the research previously undertaken uses mathematical modelling of daily price data to determine if technical analysis (TA) can produce abnormal returns (Curcio et al. 1997) (Lebaron & Lebaron 1996)(Qi & Yangru 2006)(Lucke 2000) etc. Whilst these studies give useful insight into the efficacy of TA, they are limited in use for the study of real-life trading due to both the granularity of their data (daily spot rates only) and their handling of transaction costs. Indeed Lucke "neglects transaction costs which I assume are small". A Retail FX Trader will both experience the intra-day movements of a currency pair, which might spike a currency price well beyond the final daily closing level, as well as the very real bid-ask spread costs incurred on every trade. To try and create a more realistic analysis of the effectiveness of TA, this study uses a computer simulation, using over 18 million instances of one minute "open, high, low, close" price data, covering 108 months of price movement across the four major currency pairs traded during that time period (BIS 2013). The simulation generated over 175 million trades, each of which included realistic trading costs and importantly using the software and broker limitations that a retail trader would use to conduct their trading. The aim of the simulation was to test three of the most common TA approaches; trend following, oversold / overbought levels and price breakout, as well as looking at the effectiveness of allowing your winners 'to run' whilst cutting your 'losers short' (Shefrin & Statman 1985).

2. Computer Simulation Methodology

The need for retail traders to phone their broker to place orders has long gone. "New electronic trading platforms" (King, Michael R, Carol Osler 2011) have come into existence. The platform used to run the computer simulation is Meta Trader 4, developed by MetaQuotes Software and released in 2005 (MetaQuotes 2016). MT4, as the platform is better known, is used by over 600 brokers to allow retail FX traders to open and close trades, either manually or via automated trading programs, written in MQL4, a C++ like language, known as Expert Advisors (EA) or trading 'robot'. These EAs allow trading strategies to be fully automated, with the computer program (EA) analysing live or historic price data, in order to make and execute trading decisions. Live price data for the available currency pairs is streamed real-time into the MT4 platform, with historic price data, across multiple time frames being available for back-testing purposes. By simulating the cumulative results from trading decisions made by the TA logic, encoded in the EA, an accurate picture of historic profitability can be obtained, including both realistic transaction costs as well as allowing for the effects of price movements at one minute granularity. The MT4 platform allows this back-testing to be carried out across multiple iterations, with each run altering internal parameters in a prescribed manner, for example to test for the effect of changing the size, frequency or exact entry criteria for a trade. The resultant data can then be exported and analysed. By running these tests across data from different years and currency pairs, the historic effectiveness of technical analysis approaches can be established. This is the approach take in the study.

To try and obtain results that were representative of the breadth of retail traders, the simulations were split across the 4 most traded currency pairs, namely the 'Euro Dollar' (EUR/USD), the 'Aussie Dollar' (AUD/USD), 'Cable' (GBP/USD) and the 'Dollar Yen' (USD/JPY). Additionally this was broken down into three periods of time. The first period, the four years between October 2008 and September 2011, represented a time during which the most frequently traded currency pair, the EUR/USD, opened and closed at very similar levels, which would be useful in avoiding a period where a significant, one-direction move took place. The second period, the three years between January 2006 and December 2008, spanned the heart of the recent financial crisis. The final two year period, January 2013 to December 2014, represented the latest two full calendar years of data available for the simulations. Unlike regulated stock markets, Foreign Exchange markets are subject to minimal regulation (King, Michael R, Carol Osler 2011) . There is no single 'official' price for a currency pair at any particular time, however the liquidity present in the market and the focus on arbitrage opportunity mean that any price anomaly between brokers is very minor and very short lived. For these tests, data from one broker "Forexite" (www.forexite.com 2016) has been used. The data is based on recording the open, high, low and close levels every minute, significantly more granular than the 'daily close' data used in other research discussed previously.

The aim of the computer simulation was to test three of the common approaches used by the retail FX trader, together with tests for other factors, such as the size and number of trades and the efficacy of 'holding your winners and cutting your losers'. In order to avoid the effects of optimising the tests to fit the data, each simulation was based on a simplified test as

to whether a technical condition was present and if so, a randomly timed trade would be opened, supporting the bias of the TA. For example, if the simplified test for a trend showed that the EUR/USD was trending higher, the simulation would allow randomly opened buy trades only. Additionally a second set of tests would be run where this logic is reversed, i.e. rather than the 'indicated' buy trade being taken, an equivalent sell trade would be placed instead. Additionally a third set of tests, where no technical analysis was present and trades were placed completely randomly, would be used as a control. Each set of random tests would be repeated 200 times for any given set of conditions, using a different 'random seed' and each of these set of tests repeated across the four currency pairs and three time periods. By comparing the data of the results from the 'correct TA', 'reverse TA' and 'base random' test runs, any 'better than random' results seen from the 'correct' use of TA should be clear, with the corresponding 'reverse' use producing worse-than-random results.

To run the computer simulations, a trading robot called "Dom" was written in MQL4, to generate random trades in line with the aspect of technical analysis being tested. Dom simulated a retail trader, with a starting pool of money in a broker account and measured the effect of the trading on the profit (or loss) on the account, including transaction charges. The base scenario for these tests used a standard trade size of 0.1 lots (equivalent to 10,000 base units of currency) together with a bid-ask spread of two pips. A starting account size of $100,000 was used, to minimise the risk of an account running out of funds and skewing any profitability data. In addition to each trade being placed at a random time, the stop-loss level, i.e. the price level where a losing trade would be closed, was set randomly between 10 and 50 pips away from the opening price, with the corresponding 'take-profit' level being set as a multiple of this, in line with the testing requirements. By using random entry timing and stop loss levels, we can simulate the action of a group traders using different entry and stop-loss criteria rather than creating an arbitrary non-random rule for trade entry which might itself unintentionally impact the results. Each test run aimed to generate a maximum of 500 trades, spread across the time period. Data from each set of the 200 test runs, repeated across the 4 currency pairs and 3 time frames was then combined, giving a total of 2,400 test runs per scenario, each run capped at 500 trades, giving a maximum target of 1,200,000 random trades per scenario. This increased to 12,000,000 trades per scenario where additional attributes, such as the reward to risk ratio was being varied between values of 1 and 10.

2.1. Scenarios Tested

A cursory glance at any of the on-line forums used by retail traders to discuss TA systems will tell you that there are many different systems in use, often combing multiple facets of TA. The purpose of this study is to determine if the use of TA, in general, can be shown to have any merit. Rather than testing one specific system, which might be unrepresentative of retail traders and a whole, the simulations will look at three concepts often found as part of a complete system; trend following, oversold /overbought levels and the concept of price breakout / extension. To test if these concepts might have any merit for the Retail FX Trader a simplified set of rules were created and built into the trading robot, Dom. In order to minimise the risk of over optimising the settings, all three use standard time settings based

around days, weeks or months, as well utilising default settings where standard MT4 TA functions have been used.

- **Trending** – if the daily simple moving average (SMA) is above the weekly SMA and the weekly SMA is above the monthly SMA, then the currency pair is trending higher (and conversely 'below' for a lower trending pair.
- **Breakout** – if the price of a currency pair is in the above the upper 10% of the previous five days range, it is considered to have broken higher (and conversely if below the five day lower 10% level).
- **Oversold / Overbought** – if the value of the default daily Stochastic indicator is above 80, the currency pair is considered to be 'over bought' (and conversely 'over sold' if the value is below 20).

These rules are deliberately simplified and non-optimised, however over a large number of randomly timed trades, any better-than-random insight they may give the trader should be evident. A positive outcome in terms of increasing profitability by the use of these TA elements will show that there could be merit in a trader using these elements as part of an overall trading system. In order to measure this a **base random test**, where no TA is in place, was also required. Additional tests to the base random case were run, to establish the impact they might have on trader profitability.

- **Number of trades** – testing the impact of transaction costs on performance.
- **Size of trade** – testing the impact of increasing size of trade on profitability.
- **"Letting your winners run"** – testing the impact of two methods of 'cutting your losers short and letting your winners run' – by price levels (Reward to Risk – "R/R" - ratio) and length of time a trade is held for.

2.2. Theoretical Outcome

In order to establish if there is any benefit (or 'edge') to the trader from these TA methods, we need to establish the theoretical mean outcome for a random trade in our simulation. For our base random case, we are randomly choosing whether to buy or sell, with no knowledge of price history or fundamental information about the pair we are trading. We have a 50% chance or randomly making the correct decision and given enough random trades, these chance events should cancel each other out, leaving only the transaction costs for each trade (which have a 100% probability of being applied!) to be accounted for. There are two major transaction costs that affect traders;

- The 'bid-ask spread' (known simply as the 'spread') – the difference between the buy and sell prices being offered by the broker. This spread is measured in the base unit of price movement, known as pips.

- The 'swap' – this is the interest rate differential between the two currencies that make up a currency pair. It is normally charged on trades that are carried overnight.

For our simulation, we are assuming a spread of 2 pips, which will be applied to all trades. Because our simulation is measuring outcomes in profit (or loss) we need to convert this 2-pip spread into US dollars (the trader's account currency we are using for the simulation). With a spread of two pips and a trade size of 0.1 lots (10,000 base units of currency) US Dollar denominated currency pairs (for example EUR/USD, AUD/USD and GBP/USD) every trade will have a transaction cost of $2.00 due to the spread. For non-USD denominated pairs, such as the USD/JPY the dollar spread cost will be dependent on the denominator's value at the time. For the simulation we are using a mean value across the three time periods based on a mid-point exchange rate for those periods. The final spread figure is calculated as a mean of the four currency pairs overall. For any set of data from the simulation that combines all four currency pairs and three time periods, we are assuming a transaction cost impact of $1.97 per trade.

Date	USD/JPY	2 pip spread cost
June 29th 2007	123.51	$1.619
March 31st 2010	93.45	$2.140
Dec 31st 2013	104.96	$1.905
	Mean USD/JPY	**$1.888**
	Mean simulation	***$1.972***

The MT4 platform we are using for the simulation does not support swap calculations, so in effect we are assuming a zero transaction cost. This will not affect the comparison between different simulated models but will need to be added in manually to make valid comments about how a retail trader might benefit from any 'edge' found.

3. Test Results

The results of the trader profitability simulations are set out below. Firstly we will look at the base random case and the impact varying the number of trades and the size of those trades has on profitability. We then examine the efficacy of 'cutting your losers short and letting your winners run' when applied to random trading. Finally we examine the three simplified technical analysis rules to see if they lead to greater profitability, versus random trading.

3.1. Base Random Profitability

The first simulation carried out was a random base test, where trades were placed with the time, direction, stop-loss and take-profit levels all randomly decided within the parameters discussed earlier. Across the three time periods and four currency pairs, 2,400 test runs generated 10,342,928 random trades.

Base Random Results							
Number of test runs	24,000		Greatest Profit (from single run)	$4,643.83		Profitable runs	6048
Number of trades	10,342,928		Greatest Loss	-$6,344.91		Losing runs	17952
Mean trades per test run	431		Mean Profit	-$846.88		Percent profitable	25.20%
Mean profit / Loss per trade	-$1.965						

These results give a simulated loss per trade of $1.965 which is less than half of one percent variant to the expected figure of $1.972. The mean profit per trade run was -$846.88 and only around a quarter of simulated accounts ended in profit, with a maximum profit of $4,643.83 versus a maximum loss of $6,344.91. The distribution of account profitability is shown below.
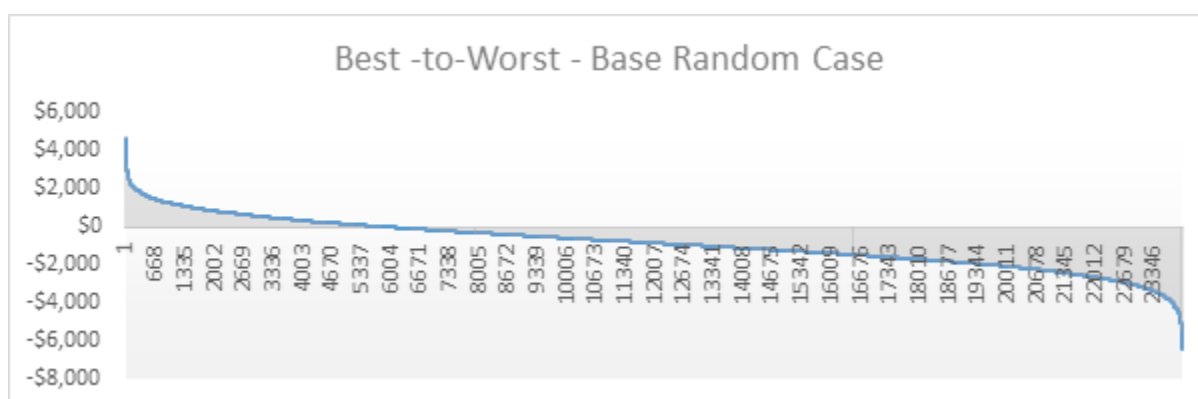
The distribution of these results, which is in-line with the author's previous work (Davison 2016) show that through random trading, whilst some traders (around 25%) will end up profitable, the mean expectation is a loss, in line with the transaction charges experienced. This is hardly a surprise, but does set the baseline for the TA simulations.

3.2. Profitability by Number of Trades

The base random trading simulation above shows that random trading produces losses in line with transaction charges, i.e. spread costs. We should therefor expect to see overall losses per simulated trader increasing with the number of trades taken in each simulation. The graph below shows the effect of increasing the number of trades from 25 to 500 in each run. The simulation was based on 2,400 iterations of random trading for each 'number or trades' case being tested. Overall 48,000 test runs generating 12,600,000 trades.

Correlation between overall profitability and number of trades
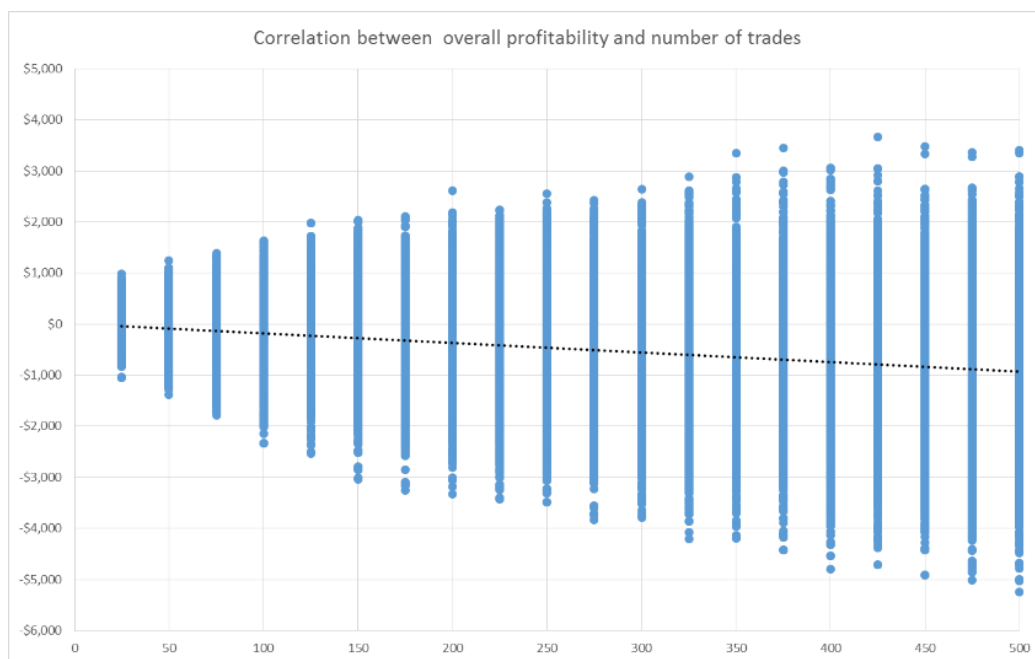
The trend line shows a mean decrease in account profitability as the number of trades increase, with the percentage of profitable accounts diminishing as the number of trades increases, from 43.4% for the 25 trade case down to 25.0% for the 500 trade sample. This is due to the constant eroding of account equity taking place as the volume of transaction charges increase. In this random trading simulation, trading volumes has an impact on both the absolute level of mean account loss and the overall chances of profitability.

3.3. Profitability by Percentage at Risk

So far the tests we have done have used a fixed trade size of 0.1 lots, or 10,000 base currency units. Depending on the position of our random stop-loss price level and the amount of funds in the trading account, the percentage of the account we are risking on any trade will be variable. Whilst our overall testing parameters have been set up to minimise the chances of running out of money in our simulated account, it is interesting to look at what happens if we do. The following tests look at the results of random trading, if we vary the trade size to risk a fixed amount on each trade. Note the trade size is now dependant on both the amount of money in the account and the stop loss level.

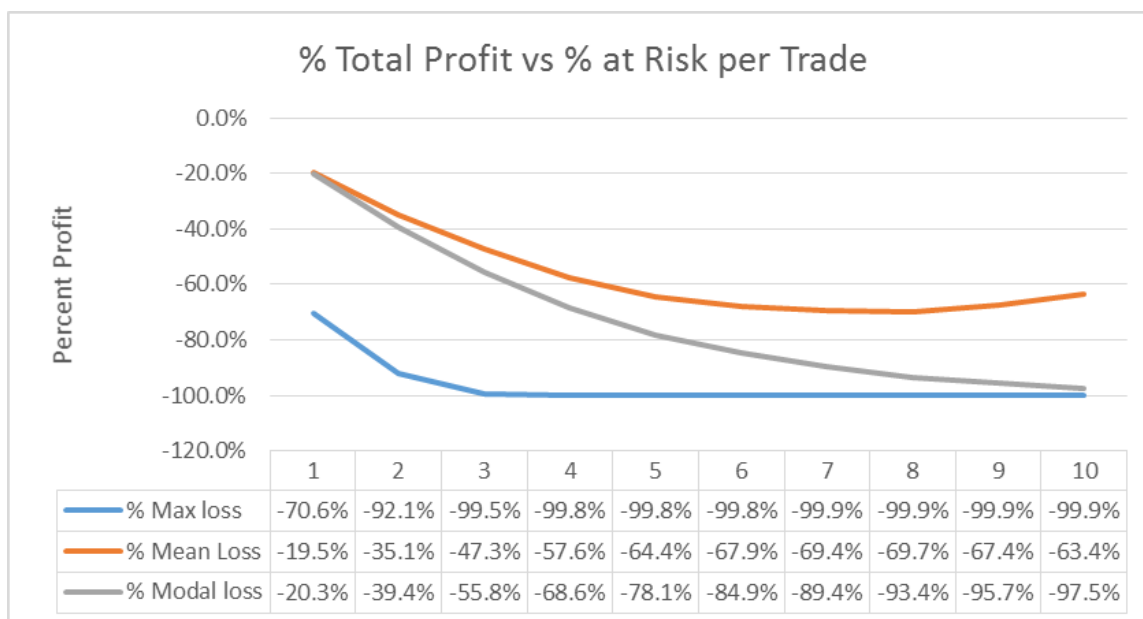| Percent at Risk - Random Results				
Number of test runs	12,000		Greatest Profit (from single run)	$1,238,554.00
Number of trades	5,033,734		Greatest Loss	($99,923.00)
Mean trades per test run	419		Mean Profit	($56,281.27)

The simulation shows that the best performing trader would have made more than a 1,200% gain, netting $1.2m by randomly placing trades, in this case risking 8% of the total account size each time. The level of this maximum gain is undoubtedly one of the attractions of retail FX trading. In comparison, the greatest loss is 'only' $99,923, however this is capped by the starting account size in the simulation being $100,000 so in effect represents total loss. Across the 12,000 test runs, only 904 (7.5%) were profitable.
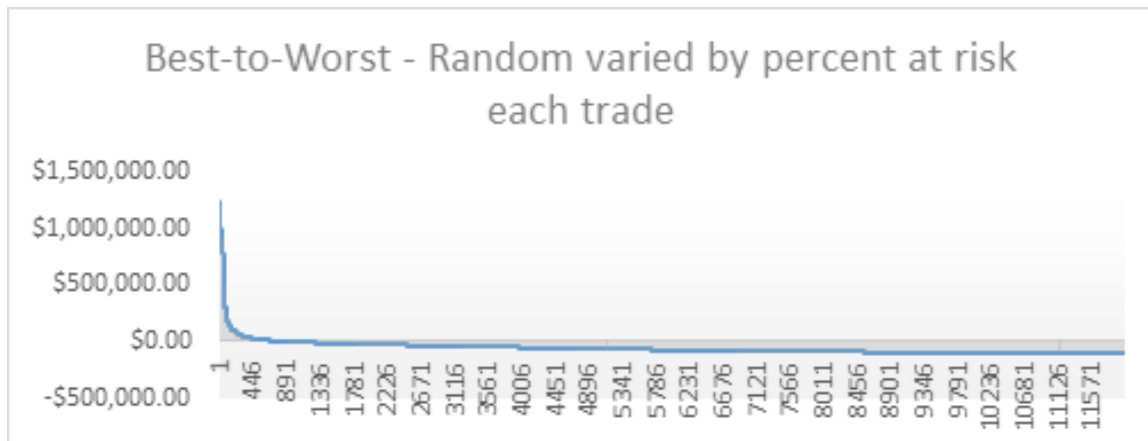
Looking at the number of profitable traders we see that in our scenario of simulating around 500 trades, risking 1% per trade has almost halved the chances of a profitable outcome from 25.1% to 12.9%. This figure is more than halved again, to 6.2%, as the percent at risk per trade increases to 5%.
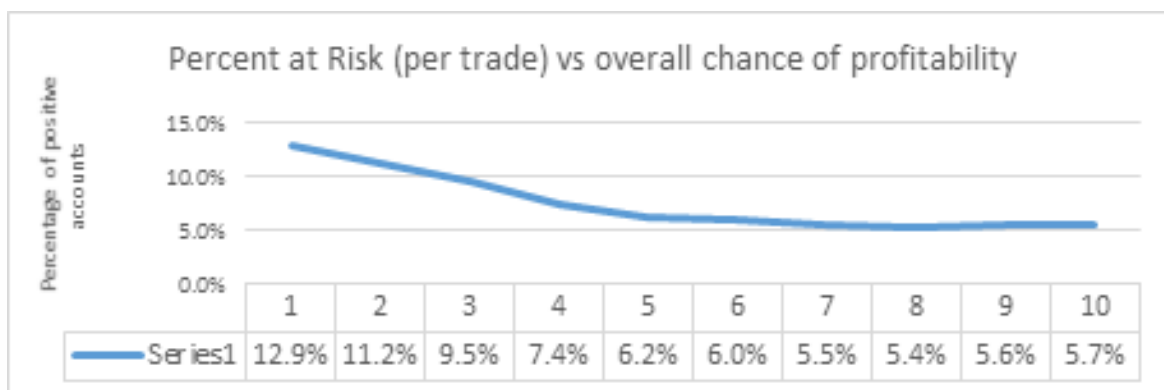
Additionally, at the 5% at risk level, the mean loss shown in the simulation was around two thirds of the total account (67.9%) with the modal loss being 78.1%. Whereas in the 'number of trades' simulation, the losses were caused by the transactional spread costs alone, in this simulation, losses are additionally created by a series of losing trades that cause trading to cease due a lack of funds.

% Total Profit vs % at Risk per Trade

	1	2	3	4	5	6	7	8	9	10
% Max loss	-70.6%	-92.1%	-99.5%	-99.8%	-99.8%	-99.8%	-99.9%	-99.9%	-99.9%	-99.9%
% Mean Loss	-19.5%	-35.1%	-47.3%	-57.6%	-64.4%	-67.9%	-69.4%	-69.7%	-67.4%	-63.4%
% Modal loss	-20.3%	-39.4%	-55.8%	-68.6%	-78.1%	-84.9%	-89.4%	-93.4%	-95.7%	-97.5%

Combining these results with that of the impact of the number of trades, we see a clear trend that shows the more you trade and the more you risk per trade, the higher the level of transaction costs you need to overcome in order to be profitable and the higher the chances of account loss that stops you trading.

3.4. Cutting your Losers Short and Letting your Winners Run

The idea that 'cutting your loser short and letting your winners run' is good practice, is not new, but it is difficult to do psychologically. Loss aversion leads to investors to "sell winners too early and ride losers too long" (Shefrin & Statman 1985). A trader who has a profitable trade may wish to close the position to both realise the profit and reaffirm their 'correct decision', whilst a losing trade may be allowed to run to allow it a chance to become profitable. To find out what impact cutting losers and letting winners run has, we can use the random trading simulator by continuing to open trades randomly, but to close them in line with one of two approaches;

- Time based – choosing a random time period over which a trade will be closed if it losing, but keeping a winning trade open for a positive multiple of that random time period. For example keeping a losing trade open for only 1 hours but allowing a winning trade to run for 5 hours.
- Price based – choosing a random price movement at which a losing trade will be terminated, whilst setting the take-profit price movement at a positive multiple of that value. For example closing a losing trade after a 10 pip movement against the trade but only closing a winning trade after 50 pips of profit. (This is known as a 'reward to risk' ratio).

Both of the simulations carried out test the effect of varying the 'positive multiple' between 1 and 10. The entry time and direction (buy or sell) are random. In the 'time' simulation, a random time to close a losing trade of between 1 and 10 hours is generated for each trade and

the corresponding time to close a winning trade set to the multiple being tested. Similarly in the 'price' scenario, a random stop-loss of between 10 and 50 pips is selected, with the take-profit level being set as a positive multiple. These maximum values of 10 hours and 50 pips have been used as they are broadly equitable without being overly optimised.
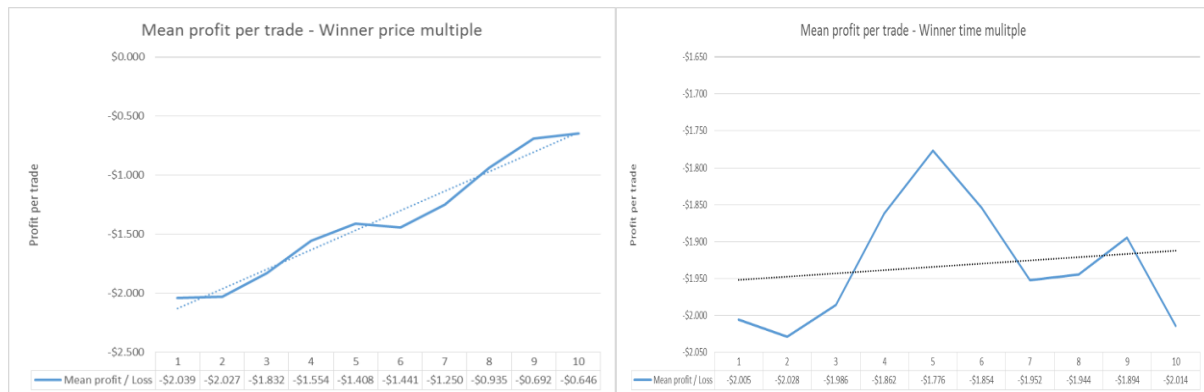

In these simulations, involving over 20 million simulated trades, there is a clear positive correlation when allowing trades to run further (price) if they are winners, however letting trades run longer (time) seems only to be effective over a relatively short period. The theory behind letting your winners run is that the value of a currency pair (or stock) might increase over a period on time as a positive trend develops. Setting a price based target, to capture any trend, will result in the trade closing when that price target is hit, regardless of the time it may take. Setting a time target in the future, will result in the trade closing at whatever the price of the currency pair (or stock) is at this fixed point in time. If the price moves significantly in a positive direction and then retraces, the price based target may capture this move, whereas the time based target might close the trade after the price has retraced. If the measure of 'success' was the time that a trade remained open, which might be the case if a currency trade was generating profits via a net interest rate payment (swap) such as might be in place for a 'carry trade' then there may be a benefit for the trader in a time based approach. For the retail FX trader however, the spreads on the interest rates used to generate any swap credit means that such net gains are likely to be insignificant to that of the price movement.
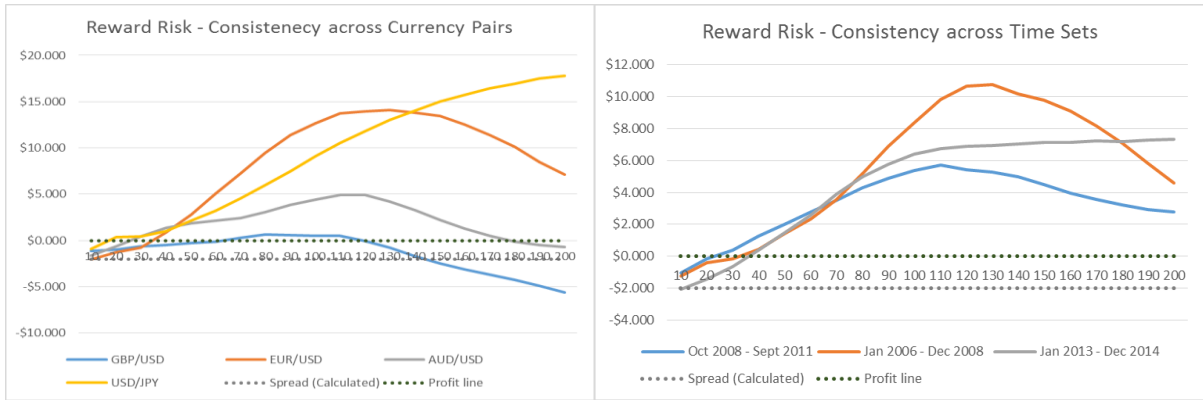
 These results are the aggregate across 3 trading periods and 4 currency pair, however breaking down the overall results of the price simulation into their constituent time and currency pair parts, also shows a similar positive correlation.
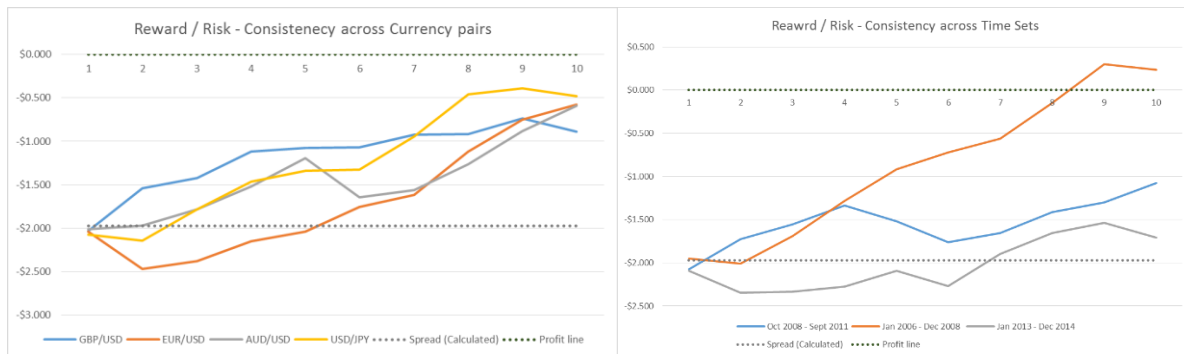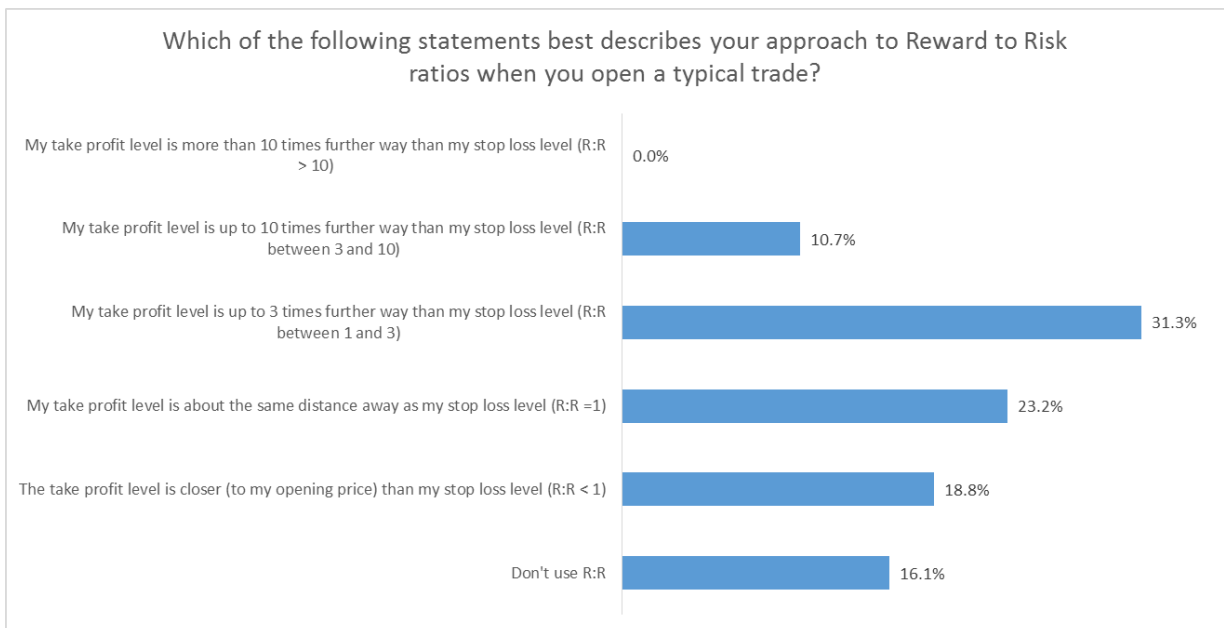
The positive results for the price scenario leads to the question of how far this positive correlation will extend? The simulation next looked at multiples between 10 and 200, again using only random trade timing and direction of trade. To try and keep the take-profit price target to a realistic number, the random stop-loss range was reduced to a value between 10 and 20 pips, meaning that the corresponding take-profit target would be between 100 and 4000 pips, depending on the random stop-loss and the multiple being tested. For comparison, a similar simulation for the time scenario was also carried out, also extending the original test to time multiples of 200.
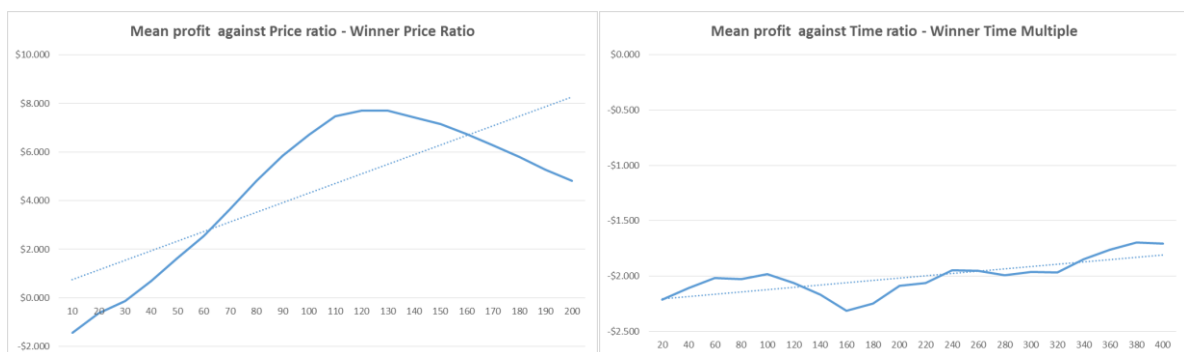

The results show once again a clear case for extending price rather than time. Whilst the time simulation shows a slight trend higher as holding times increase, with a mean profit per trade, ranging from -$2.22 to - $1.69 (ie a decreasing loss), the price scenario shows a mean profit per trade of $7.70 is being achieved at a reward to risk ratio of 130. This is remarkable in that this is being achieved without any technical or fundamental analysis, simply by setting random trades to take profit at a price level that is 130 times greater than their stop-loss level. Again this correlation is seen at the constituent currency pair and time-frame level.

Reward Risk - Consistenecy across Currency Pairs

Legend: GBP/USD, EUR/USD, AUD/USD, USD/JPY, Spread (Calculated), Profit line

Reward Risk - Consistency across Time Sets

Legend: Oct 2008 - Sept 2011, Jan 2006 - Dec 2008, Jan 2013 - Dec 2014, Spread (Calculated), Profit line

Note that the factor of '130' is intrinsically linked to the stop loss range to which it is being applied, in this simulation a random level between 10 and 20. The implication is that for these simulations, a price movement of around 1,950 pips (the mean random stop loss multiplied by the peak Reward/Risk multiple) was the optimum target using this random approach. Either side of the knee of the curve are price movements that were either sub-optimal, as the price moved further, or overly far away as the price of the currency pair never reached them. The constituent graphs show that whilst the exact position of the knee varies, in all cases it represented a reward to risk ratio of over 100 (so a mean of 1,500 pips) away and additionally, present across all 3 time periods and 4 currency pairs. This is particularly interesting in the light of the results of an online survey of Retail FX Traders, conducted by the author, in which one of the questions asked about the 'approach to reward to risk ratios'. Of 127 respondents, not one trader said they use a ratio higher than 10, and whilst there is no claim made that a higher reward to risk ratio is the only route to profitable trading, it would seem that the Retail FX Trader may be missing out by not using longer RR ratios?

Which of the following statements best describes your approach to Reward to Risk ratios when you open a typical trade?

Statement	Percentage
My take profit level is more than 10 times further way than my stop loss level (R:R > 10)	0.0%
My take profit level is up to 10 times further way than my stop loss level (R:R between 3 and 10)	10.7%
My take profit level is up to 3 times further way than my stop loss level (R:R between 1 and 3)	31.3%
My take profit level is about the same distance away as my stop loss level (R:R =1)	23.2%
The take profit level is closer (to my opening price) than my stop loss level (R:R < 1)	18.8%
Don't use R:R	16.1%

The study next looked at the case of setting profit targets closer than your loss targets, known as 'scalping'. From the survey results we see this is practiced by 18.8% of respondents, however we also see that the same correlation between RR ratio size and profitability exists.
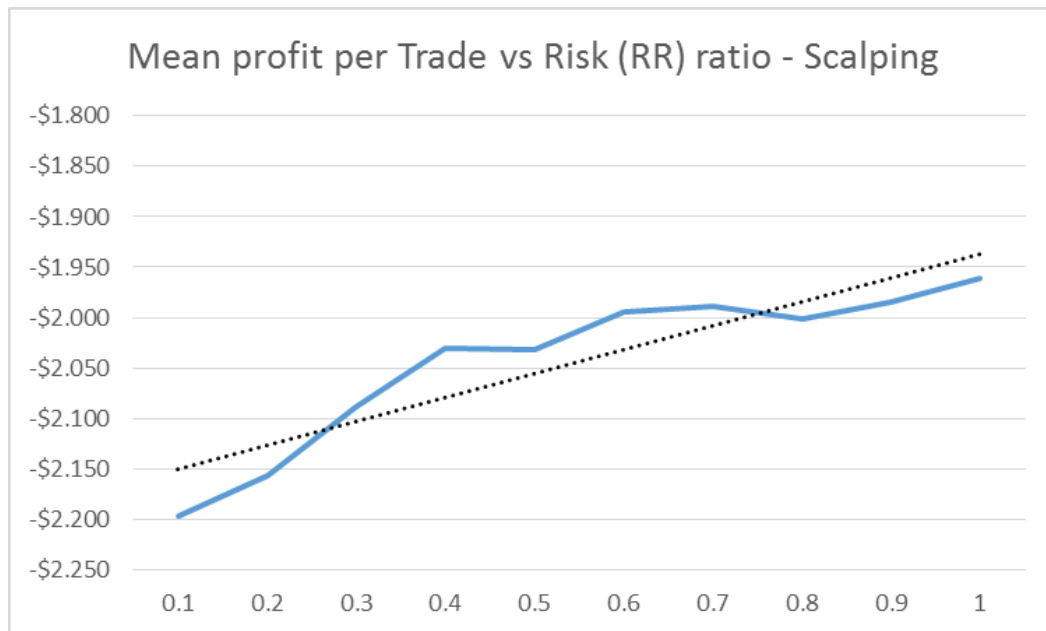

The logic of Scalping is that letting your 'losers' run may give them time to become profitable, however the simulation of 12 million random trades suggests that whilst a high win rate may be achieved, the larger net value of the (smaller number of) losing trades is not compensated for by the (larger number of) winning trades. The attraction of having more winning trades may be one of the reasons that in the survey shows over one-in-six traders trading in this way?

3.5. Technical Analysis Results

So far we have examined the impact on profitability of trades, all of which have been opened at a random time and random direction, by implication with no technical analysis (TA) taking place. As discussed earlier, TA is used by around 90% of traders and so understanding if there is any value for the retail trader is important! In the following three simulations, a random trade and direction was generated as before, however the trade was only opened if the direction of the proposed trade (i.e. buy or sell) was confirmed by the element of technical analysis being examined. In each case the stop-loss was still randomly generated with the take-profit level being a multiple of the reward to risk ratio additionally being tested. The same tests were then carried out with the signal from the technical analysis logic being reversed, i.e. placing a buy rather than a sell. In this way, positive validation of the positive TA rules will occur if it results in both higher profitability than both the inverse rules and the random bases tests carried out earlier.

Whilst the trading platform we are using to run these tests, MT4, supports realistic transaction fees in terms of the bid-ask spread, it does not simulate the interest rate contributions or charges for holding a trade overnight, known as the Swap. The Swap can either be a net credit, or debit, depending on the underlying interest rate of the currencies that make up a currency pair. For example, based on typical retail FX broker quoted interest (Oanda 2015) rates selling the (low interest rate) US Dollar to buy the (much higher interest rate) Australian Dollar (buying the AUDUSD) would results in a typical overnight credit of $0.68 per 0.1 Lot during the 2006 to 2008 period testing, whereas the inverse trade, selling the AUDUSD would cause an overnight Swap charge of $0.99.
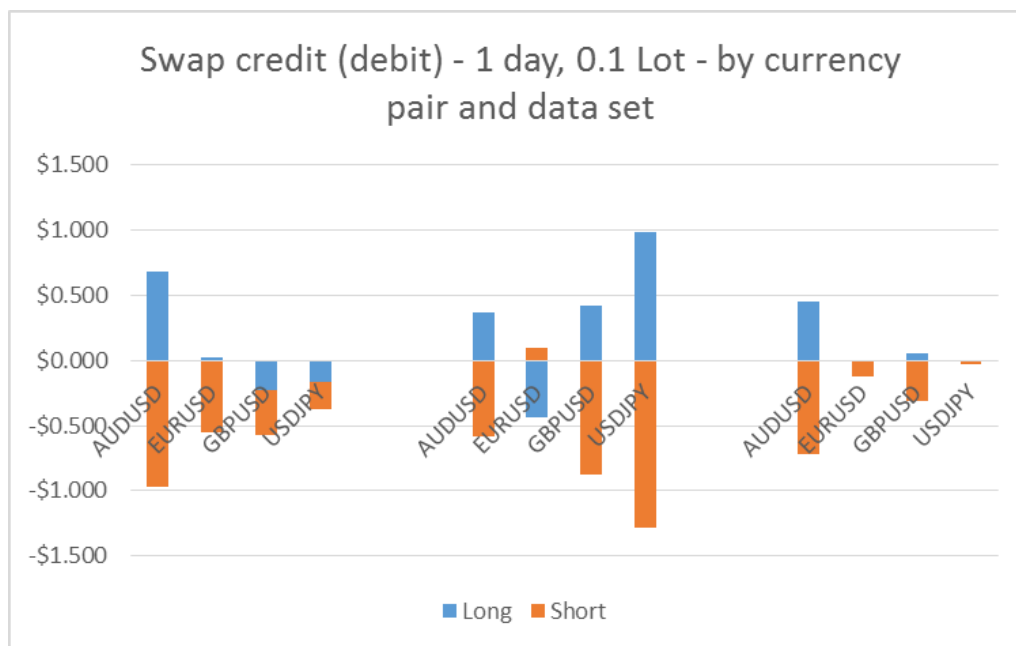

The chart above shows the mid-point overnight interest rate debits or credits associated with the three time periods used in the simulations. In order to apply the effect of the Swap to the overall results, assumptions were made as to the length of time a typical losing and winning trade would stay open (based on the RR ratio being tested). Across the simulation an equal number of long and short trades (buy and sell) were assumed as well as an equal distribution across the three time periods and for currency pairs. From this a net debit of $0.156 per 0.1 lot trade per day was calculated and factored into the overall results of each simulation run (including the base random scenario) to reflect the net effect of the Swap that a Retail FX Trader may face.

3.6. Trend Scenario:

Trend trading is based on the idea that once a currency pair (or stock) has started to move higher (or lower) over a period of time, it is likely to continue to do so. One way of testing to see if a trend is in place is to use the relative position of a series of simple moving averages of different periods. The trend simulation TA checks to see if the short term daily simple

moving average (SMA) is above the weekly SMA and the weekly SMA is above the monthly SMA, and if so will conclude that the currency pair is trending higher (and conversely 'below' for a lower trending pair). If the three moving averages are neither in descending or ascending order, then no trend is assumed.
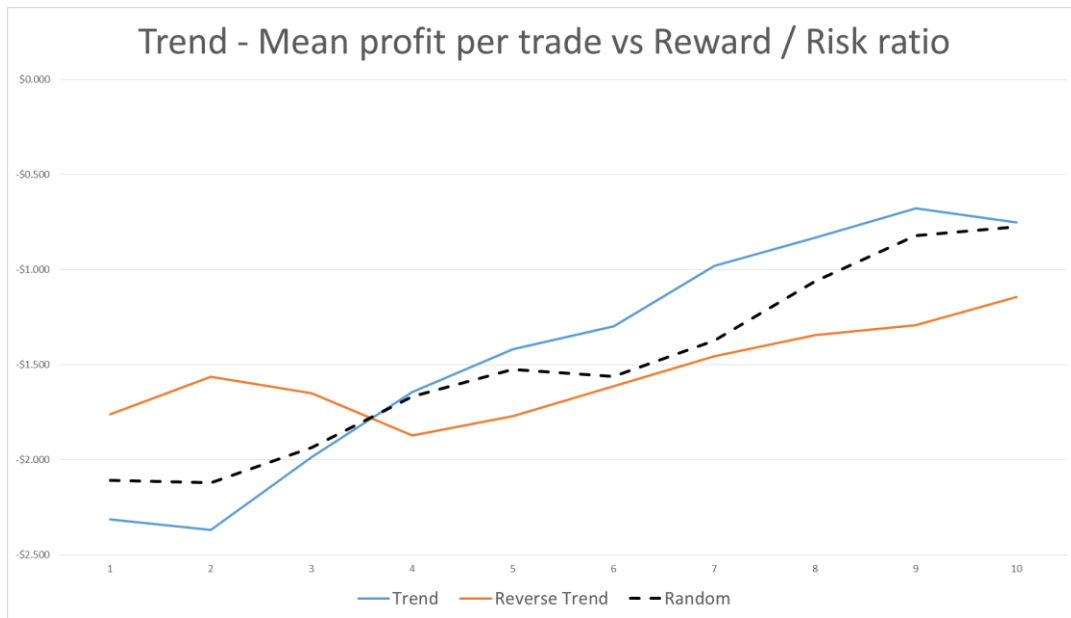

The results of the 19.5 million simulated random trades show that whilst the positive 'trend' analysis does seem to give better-than-random and better-than-reverse results at Reward / Risk (RR) ratios above 3, that there is very little real world performance benefit over random trading. This could be due to the specific simplified rules tested, or be a due to the counter trend theory, that of mean reversion, where price is viewed as being more likely to revert to a mean value, rather than to continue moving away from that mean in a 'trend' hypothesis.

3.7. Break Out:

The theory of 'break out' trading is that following a period where the price of a currency pair has been trading in a range (as opposed to moving in one direction, such as we might see in a trend) that a break-out of this range may be the start of a trend and therefor significant price move. Expanding this theory, if the price is in the upper decile of where it has been over a period of time, there may be a fundamental reason for this, which may continue into the future. For the simulation, a rule was created that said "if the price of a currency pair is in the above the upper 10% of the previous five days range, it is considered to be an upper break out (buy) opportunity and conversely if below the five day lower 10% level, a sell opportunity".
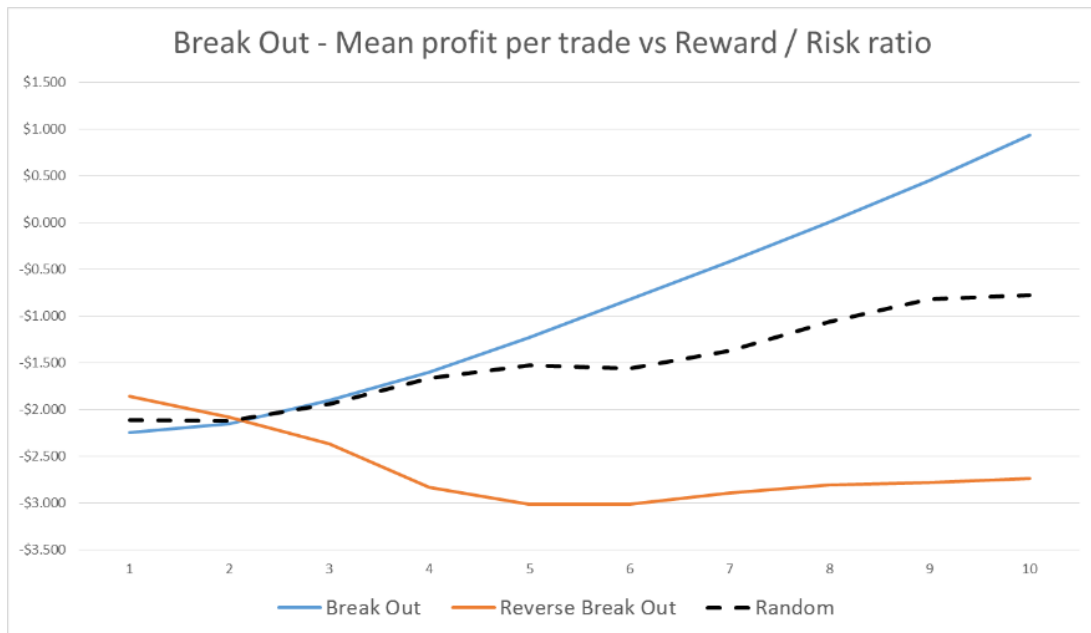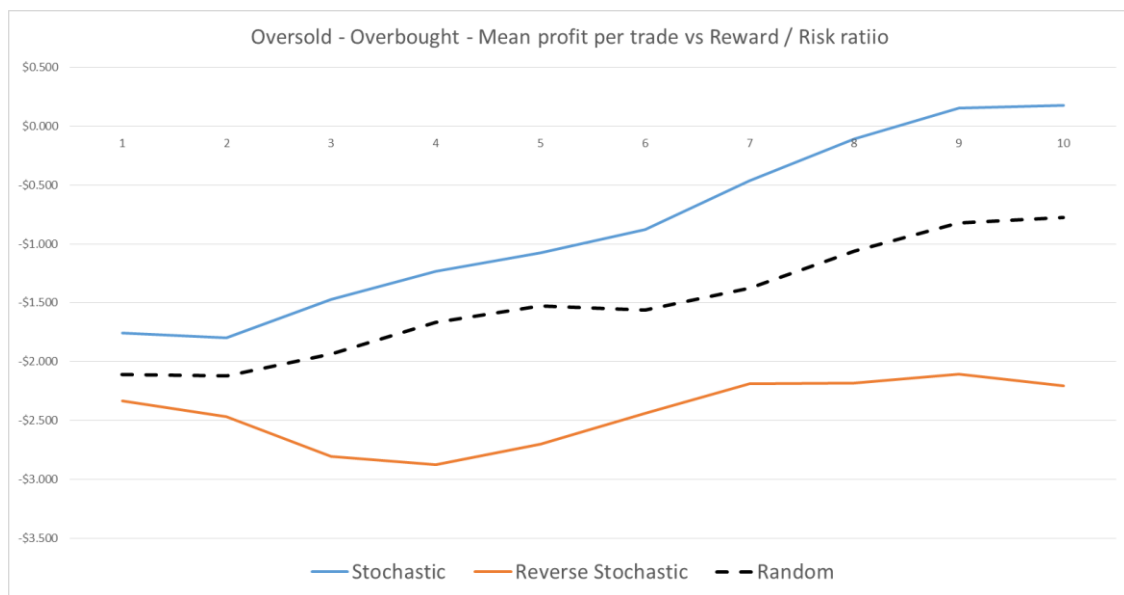
The simulation results of the positive use of this TA rule appears to show a clear profitability benefit over both the random (no-TA) results and the inverted use of the rule, for RR values greater than 2. However it is interesting to break these results into the constituent currency pair and time period data. Whilst the combined results (above) are positive, we see (below) that this hides the worse than random results that a trader would have had between October 2008 and September 2011, similarly for the GBP/USD and USD/JPY currency pairs across all three time periods.
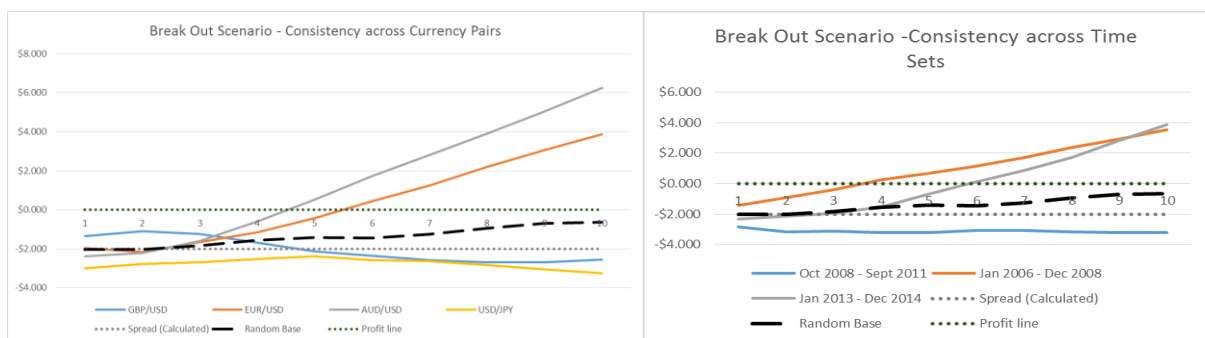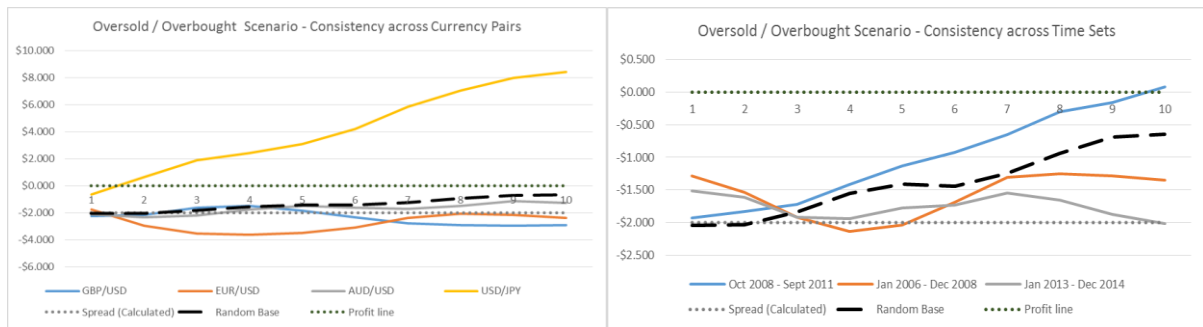

This differs from the random Reward / Risk ratio simulations discussed earlier, which shown positive results across all the data sets. The overall, net results, which were again the results of over 20 million randomly timed trades, suggest that technical analysis looking at a break out could help the retail trader achieve better than random results, although there is less consistency at the component level with the specific TA used.

3.8. Oversold / Overbought:

The final area of technical analysis for the simulation was that of the concept of a currency pair being overbought or oversold. The idea is that a currency pair (or stock) that has been subject to buying pressure over a period of time, reaches a level at which it has become 'over bought' and as such is more likely to come down in price. There are many TA tools that can be used to try and put a quantitative value on this concept however for the retail traders the 'Stochastic Oscillator' invented by George Lane (Lane 1984). Lane used a 14 day average of the ratio of daily closing price to daily range, expressed as a percentage, to try and identify changes in the momentum of a stock. The oscillator gives a value of between 0 and 100, with the simplest interpretation being that values over 80 are linked to a potential over-bought situation and under 20 indicating over sold conditions.

Using this simplified interpretation, the simulation of over 22 million trades shows a clear, mean performance benefit from the positive use of this simple TA rule. Once again however, when looking at the constituent elements, we see that majority of gains came from one currency pair, the USD/JPY and in one time period, October 2008 to September 2011. Here, profitable trading occurred at Reward to Risk values of greater than two, with a mean profit of over $6 per trade using a RR of ten. It is interesting to note that these were two of the three components that performed significantly worse-than-random in our breakout tests, which were looking for a move to continue, rather that reversing, as with these tests. One conclusion may be that these TA rules are less about predicting the movement of the price of a currency pair, rather they provide a useful entry point to capture whatever move ends up taking place?

Oversold / Overbought Scenario - Consistency across Currency Pairs

Oversold / Overbought Scenario - Consistency across Time Sets

4. Conclusion

This paper set out to answer the question 'can a Retail FX Trader ever expect to use technical analysis to achieve net profitable outcomes?' All trading has an element of risk and for every buyer there needs to be a seller, so there are no guarantees of profitability for all, however there would seem to be things that a Retail FX Trader can do to improve their chances of success. By simulating over 175 million trades, spread across the 4 most traded currency pairs and across 9 years of data, consistent patterns of behaviour that would have led to net profitable trading have been identified. This would imply that the thesis of the 'random walk' is incorrect, even if there is a random element at the level of an individual trade.

The Retail Trader needs to understand that the playing field is slanted against them in terms of the transactions costs they must pay. It would be quite possible to obtain 'better-than-random' outcomes and still make a loss due to these costs. The simulations show, perhaps unsurprisingly, that for random trading there is a clear relationship between the number of trades taken and both the amount of loss a trader may face but also the chance of them being unprofitable. These transactions costs will need to be overcome somehow in order for a trader to be profitable and as such it might be prudent for the trader to reduce trading volume (and hence transaction costs) until they have an approach that can produce better-than-random results, of sufficient 'edge' to overcome these fees. The simulations also showed (for random trading) an inverse relationship between the lot size of a given trade (ie the amount of currency being traded) and the chances of profitability. This is in part due to the risk running out of funds during an inevitable losing streak and not being able to take advantage of the inverse winning trades that may occur, especially when the win chance of an individual trade is lower. Perhaps the best chances of profitability from these two factors can be summarised as "*trade little and not-very often*" at least until a better-than-random expectation has been achieved.

The simulations also looked at the concept of 'cutting your losers and letting your winners run' and the effect of both a time based and price based, 'reward to risk ratio' approach. Whilst the traditional time based approach did show some merit, it alone did not achieve significantly better-than-random outcomes and never reached the level of showing net profitability in the simulations. By contrast, a price based approach, where the price level of closing a winning trade was a positive multiple of pips further away from the opening price than the level used to close the trade if it was a 'loser', gave significantly better-than-random and indeed profitable outcomes, even when combined with random trade entry. More remarkable was that the net profitability of this approach was consistent across all four

currency pairs and three time periods. Whilst these simulation results would seem to offer a 'holy grail' to the Retail Trader, it is worth noting that this net profitability is linked to a small number of winning trades positively offsetting a larger number of losing trades. As noted earlier, a trader will need to adjust their trade size cope with the expectation of a string of losing trades that such an approach will bring.

Finally we should consider the results obtained from the simulation testing of the simplified Technical Analysis (TA) approaches. Whilst deliberately being simplified and non-optimised, all three TA approaches showed over the course of 60 million plus simulated trades, both better-than-random and better-than-inverse net profitability.
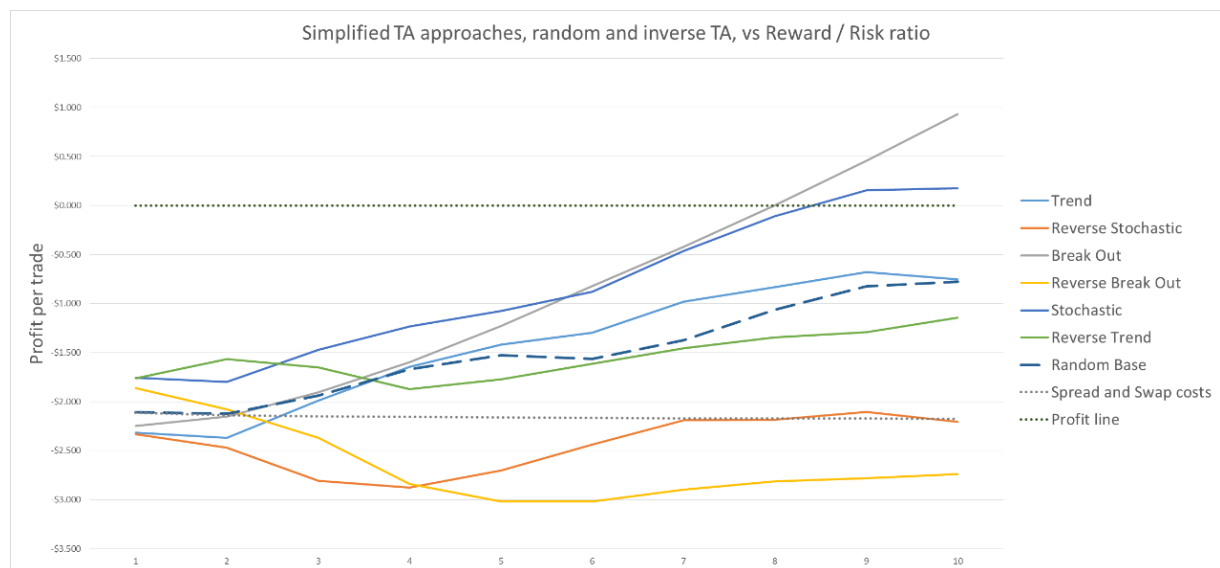
Whilst it was never the intention of this study to identify a specific, profitable trading methodology, these simulation results suggest that a TA approach has merit to offer the Retail Trader, especially when combined with a suitably large Reward to Risk ratio. No claim is made that this is the only, or indeed the best, approach for a Retail FX Trader to take, only that based on these non-optimised back testing simulations, that profitability is an obtainable goal. One note of caution however can be seen from the chart above. Whilst the use of positive Reward / Risk ratios and Technical Analysis produced net returns better than the transaction costs, indicating there is a non-random nature of the market, there is still a significant 'edge' gap that needs to be closed in order to achieve net profitability. The lower the transaction costs of a trader, the more chance they have of crossing that gap.

Of all the elements of trader profitability tested, the simulations would suggest that the use of an extended Reward to Risk ratio, based on price (rather than time) holds the key to success and profitability. This makes sense given the historic, fundamentally driven valuation changes that have affected a countries currency over recent years. The use of positive R/R levels does not require the prediction of the direction of that move, nor the timing of it, only to allow a trade to stay open long enough to capture that move and to minimise losses by

shutting down potentially losing trades early. To answer the question 'can a Retail FX Trader ever expect to use technical analysis to achieve net profitable outcomes?' the answer, based on these simulations, is a cautious 'yes'. There will always be risks and there are certainly no guarantees, however it would appear that there is hope. Trade small and set significantly bigger price targets for your winners than for your losers.

References

BIS, 2013. Triennial Central Bank Survey. Global foreign exchange market turnover in 2013. , (April), p.24.

Curcio, R. et al., 1997. Do technical trading rules generate profits? Conclusions from the intra-day foreign exchange market. *International Journal of Finance & Economics*, 2(4), pp.267–280. Available at: http://doi.wiley.com/10.1002/(SICI)1099-1158(199710)2:4<267::AID-JFE57>3.0.CO;2-J.

Davison, C.J., 2016. The Retail FX Trader : Random Trading and the Negative Sum Game . , (Jan 2016). Available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2711214.

ForexFactory, 2015. http://www.forexfactory.com/.

Heimer, R.Z. & Simon, D., 2012. The Dedicated and The Dabblers: How Social Interaction Propagates Active Investing. , pp.1–70.

King, Michael R, Carol Osler, D.R., 2011. Foreign exchange market structure, players and evolution. *Norges Bank*.

Lebaron, B. & Lebaron, B., 1996. Technical Trading Rule Pro tability and Foreign Exchange Intervention. *Exchange Organizational Behavior Teaching Journal*, 49, pp.125–143.

Levich, R.M. & Thomas, L.R., 1991. THE SIGNIFICANCE OF TECHNICAL TRADING-RULE PROFITS.

Lucke, B., 2000. Are Technical Trading Rules Profitable ? I . , pp.1–15.

Lui, Y.-H. & Mole, D., 1998. The use of fundamental and technical analyses by foreign exchange dealers: Hong Kong evidence. *Journal of International Money and Finance*, 17(3), pp.535–545.

Menkhoff, L. & Taylor, M.P., 2006. The Obstinate Passion of Foreign Exchange Professionals: Technical Analysis.

MetaQuotes, 2016. MetaQuotes. , p.http://www.metaquotes.net/en/metatrader4.

MyFXBook, 2015. Myfxbook. Available at: https://www.myfxbook.com/ [Accessed April 25, 2015].

Oanda, 2015. Forex Trading Platforms | Accounts & Practice Accounts | OANDA fxTrade Europe. Available at: http://fxtrade.oanda.co.uk/trade-forex/ [Accessed May 27, 2015].

Qi, M. & Yangru, W., 2006. Technical Trading-Rule Profitability , Data Snooping , and Reality Check : Evidence from the Foreign Exchange Market. *Journal of Money, Credit, and Banking*, 38(8), pp.2135–2158.

Shefrin, H. & Statman, M., 1985. American Finance Association The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. *Source: The Journal of Finance*, 40(3), pp.777–790. Available at: http://www.jstor.org.

Taylor, M. & Allen, H., 1992. The use of technical analysis in the foreign exchange market. *Journal of international Money and Finance*, 11(3), pp.304–314. Available at: http://www.sciencedirect.com/science/article/pii/0261560692900483.

www.forexite.com, 2016. www.forexite.com.