INSIGHT.DATA.CLARITY.

1st March 2017

ALTERNATIVE BETA MATTERS

Quarterly newsletter – Q1 2017

Introduction

Welcome to CFM's "Alternative Beta Matters" Quarterly Report.

Within this report we recap major developments of the previous quarter for Equities, Fixed Income / Credit, FX and Commodities, as well as Alternatives. All discussion is agnostic to particular approaches or techniques, and where alternative benchmark strategy results are presented, the exact methodology used is given.

We have also included one white paper and an extended academic abstract from a paper produced during the quarter. Our hope is that these publications, which convey our views on topics related to Alternative Beta that have arisen in our many discussions with clients, can be used as a reference for our readers, and can stimulate conversations on these topical issues.

Contact details

Call us +33 1 49 49 59 49

Email us cfm@cfm.fr

Quarterly review

Quantitative overview of key developments in Q4 2016

Equity indices

The final quarter of the year was full of drama with the surprise presidential election victory of Donald Trump dominating the financial and non-financial headlines. The news was initially greeted negatively by markets but sentiment quickly changed as the run up to the end of the year saw developed market equity indices benefiting from the Trumpflation trade - a fiscal deficit charged, infrastructure splurge, putting a cushion under prices and providing a potential short to medium term economic boost.

Unfortunately for emerging markets, however, Mr Trump's victory weighed heavily on asset prices as question marks arose concerning future US economic and foreign policy. The difference in performance between developed and emerging equity markets was clearly seen in the MSCI World/MSCI EM index spread - the MSCI World equity index, made up of 23 developed country equity indices, returned 1.5% through Q4, while the MSCI EM index, made up of 23 emerging countries, returned -5.6%. The correlation between the two indices continued to drift south but remained above 80%, a level still below the last peak seen in 2012. Beyond the dramatic events of the presidential election, other news affecting equity prices included the Fed raising interest rates, Eurozone equities performing well with the ECB extending its quantitative easing program and Japanese export heavy stocks rising buoyed by a weakening Yen.

Our generic trender[1] applied to equity index markets delivered good performance, with the best performer being the Eurex Eurostoxx contract. The trend applied to equity index futures was positive across all contracts, although the futures universe is strongly biased towards developed economies which all trended north in Q4. The RSI[2] applied to these same contracts reached a maximum of 64 on 21st December for the Nikkei and hovered around 50 for most contracts at the beginning of November.

The VIX index peaked at 22 points, coinciding with the result of the US election. However, this was not the high point of the year which occurred in June when the Brexit vote saw the VIX reach 25 points. Options markets in Europe and Japan showed very similar patterns throughout the quarter and, as we finish out the year, equity index implied volatility has fallen back to low levels. Liquidity levels continued to be good with risk weighted volumes rising in November and falling through the customary quiet festive period at year-end.

The returns of the MSCI World and the MSCI Emerging Markets indices for the past year

Stocks and factors

2016 proved difficult to navigate for equity market neutral investors. Our reproduction of the Fama French factors[3] exhibited a common pattern of good performance for value and poor performance for momentum. In our case, the HML factor was positive and the UMD factor was negative over the course of the year, a pattern that also applies to Q4. The period following the US election saw domestically sensitive small cap stocks outperform, leading to a good performance in the US for the SMB factor. Overall the size effect, as represented by the SMB factor, has been positive through 2016. The final quarter of the year also saw support for financials on both sides of the Atlantic, on the back of a rise in interest rates and the perception of a less burdensome regulatory environment following the Trump election. The implied volatility of energy stocks in the US dropped as OPEC (and non-OPEC) states agreed a deal to cut the supply of Crude Oil.

The Fama-French factors for the last year in Europe, Japan and US

HML Europe

HML Japan

HML US

SMB Europe

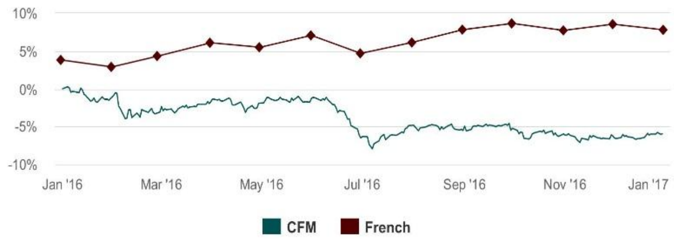

SMB Japan

SMB US

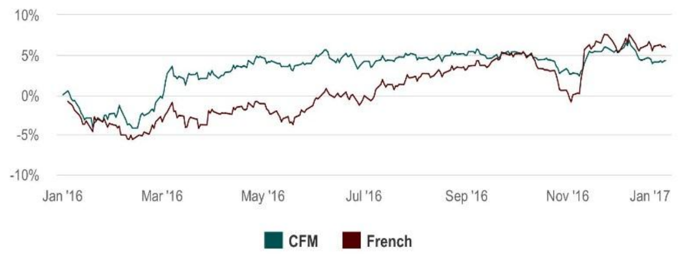

UMD Europe

UMD Japan

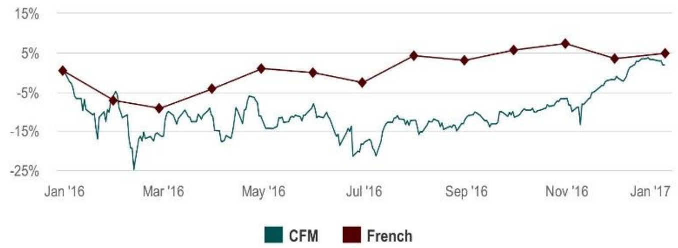

UMD US

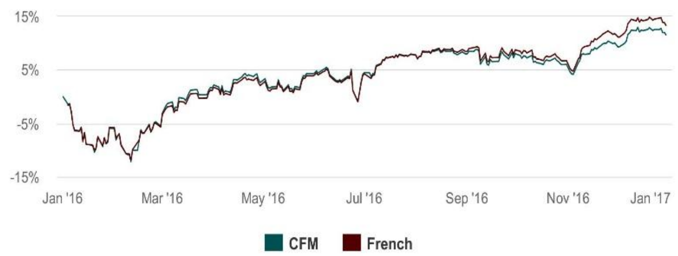

High Minus Low (HML) corresponds to a market neutral (MN) portfolio long the high book to price stocks and short the low book to price stocks. Small Minus Big (SMB) corresponds to a MN portfolio long the small market cap stocks and short the large market cap stocks. Up Minus Down (UMD) corresponds to a MN portfolio long the historical winners and short the historical losers. In each case, the red line is downloaded from Kennenth French's website, while the blue line is the CFM reproduction of the Fama-French portfolios. The methodology can be attributed to Eugene

Fama and Kenneth French and is not explicitly used in any CFM product

Fixed Income and credit

Q4 was a difficult quarter for bond holders with the Barclays Hedged Global Aggregate Bond Index reversing course and returning -2.3%. Performance on the year, nevertheless, remained robustly positive with a return of 3.9%. Bond dynamics were driven principally by political factors with the main event dominating headlines being the election of Donald Trump as US president. The Trumpflation trade, a promise of fiscal deficit fuelled infrastructure spending, was perceived by markets as providing support for prices and inevitably leading to the Fed lifting interest rates sooner rather than later. This was clearly illustrated with US 10-year yields[4] rising from 1.8% prior to the election result to 2.4% at the end of November. The CFTC's Commitments of Traders (CoT) report showed a record level of selling of non-commercial Eurodollar futures as speculators offloaded positions and moved to 2.1 million net contracts short. That same data also showed net bearish US 10-year non-commercials positioning prior to the election, hitting new bearish levels heading into 2017. In Europe, meanwhile, the ECB went to great lengths to avoid suggestions that trimming bond purchases did not constitute tapering which saw the Bund yield, at last, climbing out of negative territory.

The generic trender applied to bond markets was negative overall having been caught out in particular by the reversals in Europe. Performance was positive in the US and Japan, where yields had generally been drifting upwards anyway, while the worst bond performers were the Bund and the UK Gilts. The lowest RSI score was, perhaps unsurprisingly, seen on the US 10-year at 32 on 19th December, while the highs all hovered around the mid 50 mark prior to the US election. At the short end of the curve the reversal in the Short Sterling futures towards year-end proved costly for the trend and was the worst performer.

The US election result pushed fixed income volatilities up across the board. The TYVIX[5] index rose to 7 points in November from a low of 4 points in October, with such levels surpassing those seen in the last spike in volatility following the Brexit result in June. Risk adjusted liquidity reached the highest point of the year at the beginning of December, relaxing shortly afterwards as the markets wound down through the festive period. Liquidity conditions clearly continue to be good in the sector.

Investment grade credit generated negative returns through Q4, consistent with interest rate movements generally. High yield corporates, on the other hand, performed well, generally outperforming government bonds. In the US the BofA Merrill Lynch Investment Grade Corporate index outperformed its high yield peer by 4.8% through Q4 while in Europe the equivalent metric was 3.1%. This outperformance is also evident in the corporate CDS index market performance through the quarter.

The return of the Barclays Hedged Global Aggregate Bond and the CDX Investment Grade indices for the last year

Commodities

Q4 was a generally good quarter for commodities with the Crude heavy GSCI gaining 9.3%. Generally speaking, Q4 saw industrial commodities rise and agricultural commodities fall. Crude related instruments were lifted by news at the end of November that OPEC had agreed to a cut in production. The agreement carried extra weight following the news that OPEC had also won the backing of countries outside the cartel to reduce output, the first time such an agreement had been reached since 2001. The CoT data for Crude showed non-commercials remaining net long but unwinding ahead of the OPEC meeting and reopening long positions in December. US Natural Gas prices gained in December on the back of colder than normal weather for the time of year. CoT data showed an unwind of shorts for non-commercials with the net reaching zero at the end of December. Base metals were boosted by strong Chinese demand, with Copper, in particular, also being supported by Donald Trump's pledge of a US infrastructure splurge. Precious metals headed south, however, with the rising dollar weighing heavily especially on Gold.

The generic trender applied to commodity markets was net negative in Q4. Performance was a mixed bag with the worst performance coming from the reversals in the Crude, Brent and Natural Gas markets with Sugar and

Soybean markets also proving difficult to navigate. The trender had the most success with Cocoa's continued descent and Gold's Q4 collapse. It will perhaps come as no surprise, therefore, that the lowest RSI score came in December with Cocoa hitting 32 on the 12th of the month while the highest score was for Copper, when the red metal's impressive rally following Donald Trump's election pushed its RSI score to 69 at the end of November, prior to its subsequent fall back in value in December.

Implied volatilities generally experienced a peak in November, coinciding with the US election, and subsequently relaxed towards the end of Q4. Crude volatility, as measured by the OIV[6] index, climbed in November to a high of 55 points, only to fall prior to the OPEC meeting, falling further once the dust had settled in December. Liquidities remained good through the quarter and, once adjusted for volatility, reached a high mid-quarter before falling back through the festive period.

The one year return of the S&P GSCI

FX

Undoubtedly the FX story of Q4 was the rise in the dollar in anticipation of the Fed pushing rates up quickly to quell the potential inflation created by Trump's pledge of tax cuts and deficit spending on infrastructure projects following his surprise election as US president. The DXY, a trade weighted dollar index, climbed a hefty 7.1% through the quarter. The correlation between the DXY and the MSCI World continued to fall from being positive prior to Brexit to an approximately equal and opposite negative correlation at the end of the quarter. Q4 saw the Euro falling in response to dollar strength, while Mario Draghi successfully convinced the market that he was very dovishly tapering bond purchases, adding weight to the Euro's fall. The pound took a hit in October as Theresa May stated that Article 50 would soon be triggered to bring the UK out of the EU and looked likely to favour a hard over a soft Brexit. The CoT data showed the shortest net non-commercial GBP positioning on record at the beginning of the quarter, a trend that was reversed subsequently. The Japanese Yen also collapsed in value, a welcome change for the troublingly strong currency, bringing its value against the greenback to the level seen at the start of the year. The CoT data showed non-commercials going flat at the beginning of December and ending the quarter with new net shorts. Less developed currencies were hit hard by the election of Donald Trump which brought uncertainty in terms of future US trade policy. The Mexican Peso, a bellwether for the probability of Trump being elected, collapsed in value with the result and remained subdued heading into 2017.

The generic trender applied to a pool of currencies was net positive in Q4 with good performance coming from the European triumvirate of the Euro, the Swiss Franc and the British Pound. On the losing side was the Norwegian Krona which has been range trading all year. The Japanese Yen provided the lowest RSI score on 19th December while the most overbought currency, according to its RSI score, was the Brazilian Real, on 26th October and prior to its decline through the rest of Q4.

Implied volatility levels rose across the board in November; Euro volatility hit levels comparable to those seen in June following the Brexit result. While the single currency's volatility relaxed a little in December, demand for options saw a rally back up to similar high levels heading into year-end. British Pound implied volatility jumped in October, along with the drop in the currency itself upon news of the imminent triggering of Article 50, before falling back and remaining reasonably unchanged through to the end of the year. Risk weighted liquidity was good in the final quarter with a surge in activity in November around the US election, although this spike was smaller than at the time of the Brexit result. Liquidity conditions remain good in FX markets nonetheless.

The return of one US Dollar measured in Brazilian Real, Euro, Norwegian Krone, Japanese Yen, Mexican Peso, and the British Pound for the past year

Alternative industry performance

The final quarter of the year was difficult to navigate for the industry's biggest CTAs as the Societe Generale CTA index fell -3.8% in Q4. The year overall was not a great success, the index seeing out 2016 with a return of -2.9%. Our generic trender, once corrected for managment fees and execution costs, returned similar numbers. Average absolute correlations between the tickers in the CTA universe, an indicator of diversification, continued to fall through Q4, although they remained above the last low seen in 2015. Equity Market Neutral strategies have had a tough year in 2016 and Q4 did not buck the trend as the HFRX EMN index returned -1.2%. Among the other strategies within the HFRX database, the best performance in Q4 came from the Event Driven and Distressed Restructuring indices while the worst was the Emerging Markets index.

Total returns for Equity Market Neutral (EMN) and CTA hedge fund indices over the past year[8]

The principle implied volatility indices across four asset classes over the past year[9]

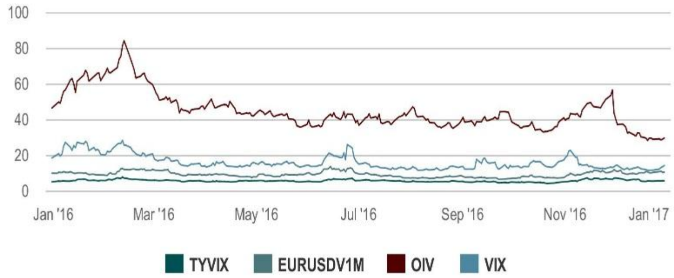

The log of the dollar risk weighted average daily volume across futures on the four asset classes over the past year[10]

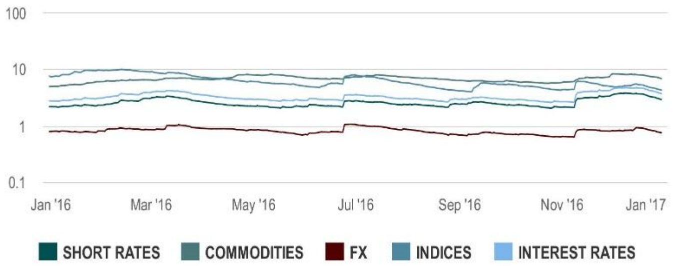

The total return of the trender defined in the text over the past year

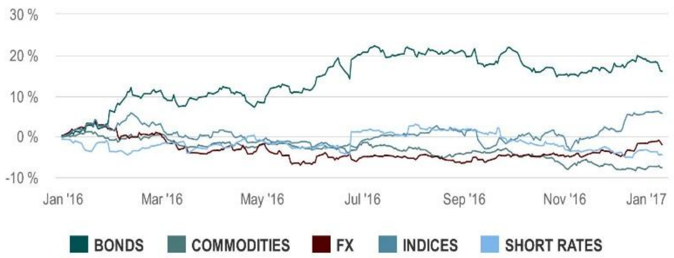

Quant of the Year award

Jean-Philippe Bouchaud is named winner of Risk's Quant of the Year award 2017

We are delighted to announce that our Chairman and Head of Research, Jean-Philippe Bouchaud, was recently awarded Quant of the Year at the Risk Awards 2017 awards ceremony in London.

The prestigious Risk Awards recognise the work of firms and individuals in derivatives markets and risk management. Judged by the editors and journalists of Risk.net and Risk magazine, they are the longest-running awards of their kind. In a line-up of 25 or so Risk Award categories, the Quant of the Year award is unique in that candidates do not enter themselves for nomination, and the winner is decided through a blend of peer and industry feedback, with the ultimate decision made by the Risk judging panel. Jean-Philippe was surprised and pleased to be declared the winner.

Risk's quantitative finance editor, Mauro Cesa, a member of the judging panel, said, "The decision to give this prestigious award to Jean-Philippe, amongst a strong field, was strongly supported by many of his peers, and his and CFM's research continues to have an impact on financial literature and market practice. We at Risk have enjoyed working with CFM over the years and look forward to their next insights."

The judging panel were impressed by Jean-Philippe's highly influential use of theory and empirical analysis in his research into finance theory. In conjunction with his role at CFM, Jean-Philippe is a physicist who has written approximately 400 papers on physics and finance where his strong focus on empirical analysis often counters traditional financial concepts. One such paper is 'Tail risk premiums versus pure alpha', written by Jean-Philippe and published in April last year, which particularly impressed the judges on this year's panel.

Jean-Philippe's current and future projects will attempt to understand the market from the bottom up. This includes market microstructure, behavioural anomalies and the impact of trading one asset on another asset's price, called cross-impact. We think this work is important both for CFM and our funds but also of significant value to the broader financial community and we're delighted Jean-Philippe has been recognised for his ongoing contributions.

You can read more about the award and Jean-Philippe's work here.

Extended abstract

Why have asset price properties changed so little in 2000 years?

Paper by Jean-Philippe Bouchaud & Damien Challet

According to the efficient market hypothesis, current asset prices should be unbiased estimators of their economic fundamentals. As a consequence, no trading strategy may yield statistically abnormal profits based on public information. The alternative hypothesis is that financial markets are intrinsically and chronically unstable. Accordingly, the interactions between traders and prices inevitably lead to price biases, speculative bubbles and instabilities that originate from feedback loops. This would go a long way to explaining market crises, both fast (liquidity crises, flash crashes) and slow (bubbles).

The debate about the real nature of financial markets is of fundamental importance. The efficient market hypothesis is not only intellectually enticing, but also very reassuring for individual investors, who can buy stocks and shares without risking being outsmarted by more savvy investors. However, a bevy of asset price anomalies have been documented in the economic literature since the 1980s, among others:

1. The Momentum Puzzle: price returns are persistent, i.e., past positive (negative) returns predict future positive (negative) returns. Trend following strategies have been successful on all asset classes for a very long time, in blatant contradiction with the efficient market hypothesis.

2. The Excess Volatility Puzzle: asset price volatility is much larger than that of fundamental quantities, and only a small fraction of price jumps can be related to exogenous news.

These effects are not compatible with the efficient market hypothesis and suggest that financial market dynamics are strongly influenced by other factors outside of fundamental quantities, in particular endogenous, self-referencing feedback loops with small perturbations potentially causing very large price changes.

A very large body of academic papers report on behavioral biases that underlie the above stylized facts. Among many others: (i) humans are strongly influenced by the behavior of others (herding) and by past trends; (ii) humans react differently to gains and to losses and prefer positively skewed to negatively skewed returns; (iii) humans are overconfident, which leads to an excess of trading activity; (iv) humans are slow to adapt to new information.

Some of these biases can now be investigated through the emerging field of “neurofinance” studying the neuronal process involved in investment decisions. One of the most salient results is that, as expected, human beings spontaneously prefer to follow perceived past trends. Various hormones play a central role in the dynamics of risk perception and reward seeking, which are major sources of positive and negative feedback loops. Hormone secretion by the body also modifies the strength of feedback loops dynamically, and feedback loops interact between themselves. Some hormones have a feel good effect, while others reinforce risk aversion. The way various brain areas are activated during the successive phases of speculative bubbles can be investigated in detail. In particular, regrets or a “fear of missing out” lead to trend following.

Human brains have most probably changed very little for the last two thousand years. This means that the neurological mechanisms responsible for the propensity to invest in bubbles are likely to influence the behavior of human investors for as long as they will be allowed to trade.

Whitepaper

Executing with Impact - why the price you want is not the price you get

Executive Summary

CFM has been trading in the world's most developed financial markets since 1991 and, since 2002, has been using in-house developed algorithms to execute through brokers and exchanges, interacting directly with electronic order books. The firm has, in this time, dedicated significant resources to understanding the microstructure of markets, publishing extensively on the subjects of slippage and, in particular, market impact. In this short note we share some of our insight and experience with references to our experimental data in order to explain the origin of trading cost, moving from bid-offer spreads to the ideas of price impact before describing a recipe for modelling these costs in simulation.

Introduction

Financial markets provide an interesting laboratory in which to study the real time dynamics of supply and demand between buyers and sellers. The most developed markets traded by CFM are generally "central limit order book" markets that are transparent, anonymous and often some of the most liquid in the world. Our experience with these markets is mainly focused on equities, futures, options, FX[11] and certain liquid bond markets. As a firm we collect every trade and order book change that has gone through these markets resulting in a data-base that grows at the rate of approximately 300Gb per day with the ability to collect such vast amounts of data arising from a significant investment in IT infrastructure. As traders of these markets we have also collected a data-base of our own trades. This data is, of course, not something that can be bought and provides us with insight into the cost of trading given the changing market conditions and the strength of our price forecasts. This then allows us to be able to precisely model the cost of trading for use in strategy simulations.

The stereotypical financial market was one in which colourfully jacketed traders stood on the pit floor and screamed orders at one another. A client wishing to trade would phone through to a broker who would then phone through to the floor for a quote. These quote driven markets still dominate for certain instruments although the quotes are more and more generated by automated electronic systems with the pits and pit traders becoming a dying breed. In these quote driven markets clients request prices from brokers who then quote a bid (price at which they will be willing to buy) and an ask (price at which they will be willing to sell). The client then decides whether to execute the trade, execute with another broker or trade at another more opportune moment in the day. Quote driven markets still dominate in the world of FX, interest rate swaps, CDS indices and many other instruments.

Financial markets have evolved in technology and transparency and in this note we will focus on those traded through a purely electronic order book. In these order driven markets, orders placed in a market show participant interest in wanting to buy or sell at a given price. One such order book is shown in Figure 1 where the bids show market interest in buying and the asks show market interest for participants who are willing to sell but necessarily at a higher price than the best bid. If a participant wants to buy/sell now, then they need to hit the ask/bid and pay a high/low price. If a participant wants to buy/sell but is more patient, they may choose to join those on the bid/ask and wait for someone to hit them at a low/high better price. These actions occur within milliseconds of each other on the most liquid markets and provide a rich source of trade data.

In this short note we start by describing where the cost of trading comes from and try to convey that the cost is not only due to the spread between the bid and the ask - in the case of a sizeable trade, participants may try to break up the trade into smaller packets in order to trade more cheaply. We next present our own data showing that the bigger the trade is relative to the volume in the market, the more expensive the trade becomes to execute. This is consistent with the idea of a big trade pushing the price to ever more expensive levels, showing how this price impact begins to be a bigger part of the cost than spread costs. We then address the issue of brokers guaranteeing the close, which does not mean one trades for free, before concluding and directing the reader to an extensive list of CFM's academic papers.

Figure 1 – An example of a fictitious Central Limit Order Book (CLOB). The bids (buyers) are on the left while the asks (sellers) are on the right. An impatient buyer/seller will have to pay the spread with a market order (described in the text), hitting the lowest ask/highest bid and incurring a cost relative to the fair mid-price. A more patient buyer/seller could send a limit order (described in the text) at the highest bid/lowest ask price and wait for someone to cross the spread. This trade will, however, be at the back of the queue and executed on a first-in-first-out (FIFO) basis

The cost of trading – from bid offer spreads and commissions to impacts

Confronted with an order book or a broker's bid and offer, an investor naturally only considers these currently available prices when trying to estimate the cost of executing his trade. Considering the order book in Figure 1 the cost of execution comes from the spread or the difference between the best bid at 100 dollars and the best ask at 101 dollars. Assuming the price does not move, remaining static for the time of our trade, we can buy one share at 101 dollars and sell it back at 100 dollars with a loss of 1 dollar for 2 shares traded, making a cost per share of 50 cents. We could also consider the cost of trading either share as being the cost relative to the fairly priced mid-point at 100.5 dollars, which again makes each trade cost 50 cents for 1 share. Commissions are even more easily accounted for, simply being added to this per share cost. For example, for commissions of 50 cents per share then the total cost for trading one share would simply be 50+50 cents or 1 dollar.

Unfortunately, evaluating costs is never as simple when we consider a trade that needs to take more than the typical volumes available at either the bid or the ask. In the above configuration, for example, the volume available at the bid and ask may be 500 shares on each side. We may have a 100 000 share buy trade which then requires sequential trades to be placed on the market in the form of market orders hitting the ask or limit orders patiently waiting to be hit at the bid. We can evaluate the cost of each trade at the point at which they were executed relative to the fairly priced mid-point but this does not account for the fact that, due to our trade flow, we may be pushing the market up, making subsequent trades more and more expensive relative to the price before we went on the market. It is this price impact, a real-life example of which is illustrated in Figure 2, that is so problematic in the understanding and evaluation of costs. Indeed, without such price pressure occurring then life would be too easy and all strategies infinitely scalable, as the cost of trading would remain unchanged as the size of the trade increased. This does not seem plausible and, as any serious investor knows, all strategies are capacity limited due to an increase in cost with the size of the trades.

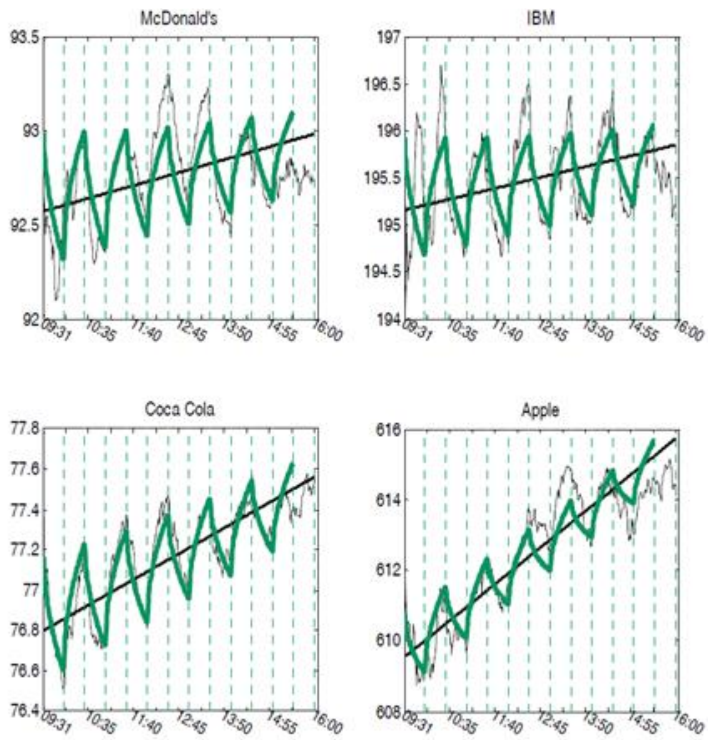

Figure 2 – The price evolution of four stocks on the 19 July 2012. This oscillating, saw-tooth behavior is more than likely caused by option hedgers in the market. Hedging options that are close to expiry and restricted to a small number of strikes (or even only one) can generate large sequential trades in opposing directions. The impact of these trades on the price of the stocks is striking and is well modelled by Equation 1. Source : Cheuvreux Quantitative Research

A trade which is split up in order to be executed over a time window forces us to think of cost as a statistical quantity. With the previous discussion of a static order book and buying and selling instantaneously, the cost remains the same no matter how many times the exercise is repeated. Breaking up the trade and executing slowly, however, means that each scenario is different and the measurement of cost now becomes an exercise in averaging over trades. Having a precise estimate of one's cost now requires many trades to be analysed and, in evaluating the worth of an algorithm or a broker, an individual trade is quite meaningless.

We now need a definition of cost in this framework of one meta-trade as being the sum of many small trades executed through a given period of time. The definition we use at CFM is the implementation shortfall, first introduced by Perold in 1988[12]. We assume that the meta-trade has been fully executed and that it is unpolluted from the impact of prior trades[13]. The implementation shortfall is a measure of the difference between the average executed price and the price before any trades were executed, as illustrated in Figure 3. This measurement of costs with meta-trades (meaning constructed with many individual trades) is, as previously described, a statistical variable. For example, executing trades over a whole day requires a data set of 100 000 meta-trades to measure an average cost precise to 1bp[14] for typical stocks.

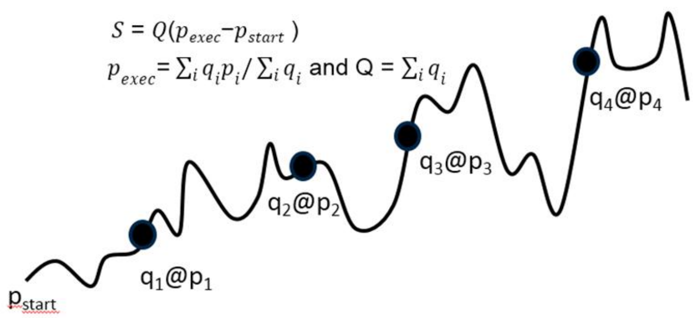

Figure 3 – An illustration of implementation shortfall, the measure of costs used by CFM. A meta-trade of Q shares is split up into 4 trades of q_1, q_2, q_3 and q_4 executed at p_1, p_2, p_3 and p_4 respectively. The cost of each can be evaluated relative to the price p_{start}, before going on the market. The total cost of the trade is conveniently expressed as the difference between the quantity-weighted[15] average price p_{exec}, and the price before trading p_{start}.

Modelling transaction costs - impact naturally explains the capacity constraint of strategies

Research in finding trading strategies begins first with trying to find a Profit and Loss curve that rises with a statistically significant level of performance. An equally important part of this research process involves trying to estimate how much the strategy will incur in trading costs once ported from a paper traded model through to real-life trading. Modelling costs is essential for this and CFM's execution research team is responsible for building these cost models. We construct estimates of costs that are robust to all market environments which is an improvement compared to measuring average costs over all recorded trades and applying that cost blindly to a strategy simulation.

If we measure the average cost of executing a stock at 5bp, for example, we could simulate with that constant cost and have a fairly stable environment in which to back-test strategies. One could clearly improve the cost model, however, by accounting for relevant changes in the market environment. The fact that costs increase in proportion to volatility, for example, seems plausible. If volatility increases then the average execution price of a buy meta-trade should increase and diverge more from the starting price. It seems reasonable that if uncertainty regarding a stock or instrument is high then the market's response to a trade, in one direction or another will be high and, indeed, higher than for a stock with a certain future. One can then build a slightly better model of costs as:

Cost per share C = a fixed (average) fraction of daily volatility

This is a superior explainer of costs but still fails to explain how they evolve as the size of the program increases. It has to be the case that as a strategy is made to manage more money that costs should go up and strategy performance down in order to constrain capacity. Impact serves this purpose! As the size of a trade increases then the price pressure also increases and the average price paid diverges from the initial price prior to going on the market, thus increasing the cost of the trade. We therefore introduce an improved cost model again as:

$$C \propto \sigma \sqrt{\frac{Q}{V}}$$

Equation 1

where σ is the volatility of the day, Q is the quantity in the meta-trade and V is the total volume traded that day. This seems to be a universal law which has been referred to in the academic and broker community, including by CFM [1]. This has also been confirmed on many asset classes such as options [3], OTC trades [4] and even bitcoins [2]. This law has been very stable in time, through different volatility regimes and, perhaps surprisingly, unchanging with the advent of a market more dominated by so called High Frequency Trader (HFT) liquidity providers. The existence of the volume term V in Equation 1 also clearly explains the motivation for adding new sources of liquidity to a universe of traded instruments as that extra liquidity increases V and thus decreases mechanically the overall cost of trading.

The existence of this square root rule is curious indeed. One observes surprisingly, for example, that 1+1=√2≠2, in the sense that two sequential trades (in one meta-trade) do not generate twice the impact of a single trade[16]. For any given strategy, increasing the capital allocated (Assets Under Management (AUM)) increases the cost incurred (unsurprisingly!). The rule in Equation 1 now tells us by how much - a doubling of AUMs leads to an increase in costs of √2=1.4 and an increase of AUMs by a factor of four leads to an increase in costs of √4=2. This point is shown clearly in Figure 4. Also of note is that as a meta-trade is executed, the price through the trade will also evolve as a square root. This is illustrated in Figure 5 below and shows that the impact of the second half of a meta-trade is much less (to the tune of about 60%) than the first. As the trade progresses the price increases, resulting in more participants being interested in selling, that generates resistance to the trade and a lessening of impact! A final surprising observation concerns the relative impact for small trades compared to big trades. A small trade generates an anomalously large amount of impact – for example, trading 1% of the average daily volume impacts the price by 10% of the volatility while trading 10% of the average daily volume impacts the price by only 3 times more or 30% of the volatility!

These observations lead us to conclude that Equation 1 dictates how much capacity a strategy has. In Figure 4 we see that as AUMs increase, the performance of the strategy decreases and at some point the strategy becomes flat and negative as the gains get eaten up by costs. This modelling of impact is therefore crucial to knowing how much a manager can allocate to a strategy.

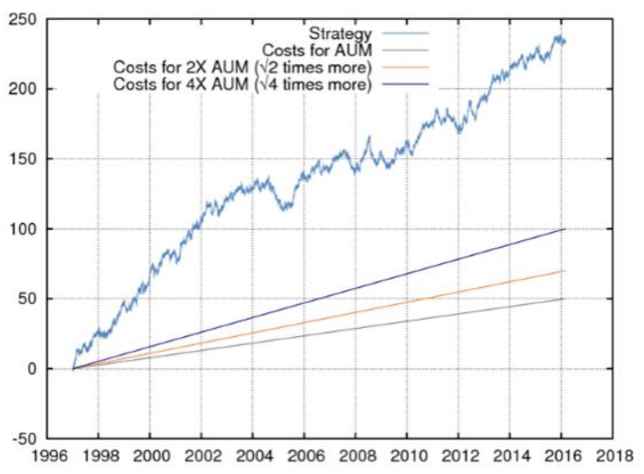

Figure 4 – The "Profit and Loss" (P&L) curve for a fictitious trading strategy with a level of costs corresponding to a given level of Assets Under Management (AUM) allocated to the strategy. Also plotted is the increase in costs following a doubling of AUM, showing an increase in costs of √2=1.4, along with the costs following an increase of AUM of a factor of four, in which case costs increase by √4=2. The P&L itself is always normalised to have the same risk for each level of AUM

My brokers are guaranteeing me the close price. Am I trading for free?

It has become standard practice in the broker community to execute client orders with a guaranteed close price[17] and in the absence of an understanding of impact this may seem like a good deal. Figure 5 below shows how a particular scenario may play out in the presence of price impact with the client guaranteed the close price at the start of the day. Assuming a constant traded volume through the day then the price will, on average, evolve as a square root, rising quickly at first and then less and less (but nonetheless continuing to rise!) as the day evolves. The average of all executed prices generates an average price of 2/3 of the total daily move (the average of a square root function which can be used in the implementation shortfall definition of cost) meaning that the broker buys at a price of p_{close}-1/3(p_{close} - p_{open}) and sells to the client at a price of p_{close}, thus pocketing the difference 1/3(p_{close}-p_{open}), as always, on average. This profit, interestingly, increases as the volatility of a market increases! A client may be reassured in volatile markets that they have a guaranteed price but the broker's profit is actually increasing with volatility!

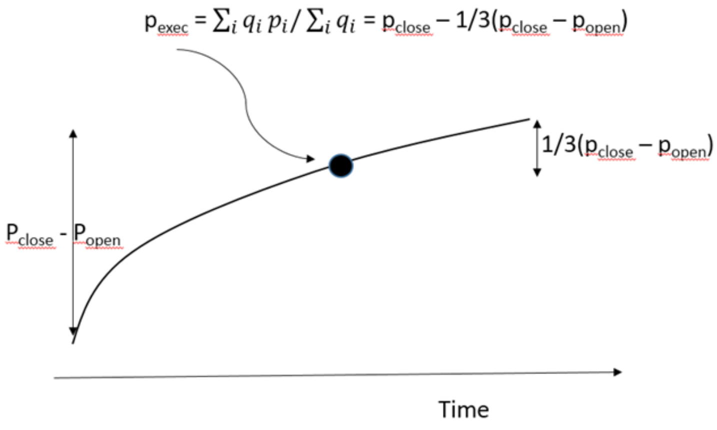

Figure 5 – The average price evolution of a trade executed from the open to the close of a market day. The black dot represents the average execution price with the quantity acquired subsequently sold to the client at the close price. The broker pockets the difference between the two prices which represents the real cost to the client

Guaranteeing the close is an innocuous practice if all volume is concentrated at the close of the market in an auction. In such a case market participants cannot impact prices through trading as no trades are executed until the end of the auction, all trades generally instantaneously occurring at a unique price and matching the maximum number of buyers and sellers. Conversely, if the close of the market is illiquid then the broker can trade through the very liquid day and then at the close of trading, push the close price as high as possible, the illiquidity meaning that small trades have large impact (we are not suggesting that brokers do this!), thus maximising his profit.

Certain market participants may, of course, need to guarantee the close, such as those hedging positions, in which case asking a broker for the close price can make sense. However, if the order is not too sizeable it may be better to just trade directly at the close through a broker rather than telling him at the open that the close price is needed. The key takeaway point is that a guaranteed price, be it the close or any other benchmark, does not necessarily mean a lower cost for your trading.

Conclusions

We have described the difference between trading costs when purely considering bid-ask spreads in an order book and those arising from price impact. The extension of these ideas for measuring and modelling costs to a trade where one is forced to split the meta-trade into small chunks is non-trivial. This transition to the world of meta-trades then requires a different technique for estimating costs and necessitates a leap to thinking in statistical terms – it is very difficult to evaluate the quality of execution based on a small number of trades! CFM has been researching execution for nearly 15 years with an aim to reducing, controlling and modelling costs. Research continues in this area with ongoing effort in the direction of better estimating costs; potentially reducing our impact or "footprint" in the market; understanding the impact of others to generate trading model ideas; and cost modelling improvements that help with portfolio construction techniques. Beyond the scope of this text is the interesting and important subject of what happens after a trade – does one's impact decay or stay constant, which has consequences for costs. Cross impact, or whether trading Apple stocks impacts Microsoft for example, is also of utmost importance, in particular for market neutral portfolios, making the cost of trading lower in such cases. These subjects, and others, will doubtless be the subject of future explanatory white papers.

References

Our references can be split into the following categories:

- Empirical analysis of trades and meta-trade data

Showing impact increases as a square root as a function of size in [1]. The square root also holds for options [3], and bitcoins [2]. Market impact slowly decays [5]. More recently we have worked on cross impact [6] and impact in OTC markets [4].

- Agent models for market impact

It is quite challenging to build a mathematical model that reproduces a square root impact. Following the model developed in [1] that was based on the idea of latent liquidity in the order book, various other models were developed, for example [7].

- Optimal trading with costs

A challenging mathematical problem concerns how to trade optimally in the face of transaction costs. We have published two papers on the subject ([8] and [9]).

1. Anomalous Price Impact and the Critical Nature of Liquidity in Financial Markets. Bence Toth, Yves Lemperiere, Cyril Deremble, Joachim de Lataillade, Julien Kockelkoren, Jean-Philippe Bouchaud
2. A Million Metaorder Analysis of Market Impact on the Bitcoin. Jonathan Donier, Julius Bonart
3. The Square-Root Impact Law also Holds for Option Markets. Bence Toth, Zoltan Eisler, Jean-Philippe Bouchaud
4. Price Impact without Order Book: A Study of the OTC Credit Index Market. Zoltan Eisler, Jean-Philippe Bouchaud
5. Slow Decay of Impact in Equity Markets. X. Brokmann, E. Serie, J. Kockelkoren, J.-P. Bouchaud
6. Dissecting Cross-Impact on Stock Markets: An empirical analysis. Michael Benzaquen, Iacopo Mastromatteo, Zoltan Eisler, Jean-Philippe Bouchaud
7. A Fully Consistent, Minimal Model for non-Linear Market Impact. Jonathan Donier, Julius Bonart, Iacopo Mastromatteo, Jean-Philippe Bouchaud
8. Optimal Trading with Linear Costs. Joachim de Lataillade, Cyril Deremble, Marc Potters, Jean-Philippe Bouchaud
9. Optimal Trading with Linear and (small) Non-Linear Costs. A. Rej, R. Benichou, J. de Lataillade, G. Zérah, J.-Ph. Bouchaud

Important Disclosures for the Whitepaper

ANY DESCRIPTION OR INFORMATION INVOLVING INVESTMENT PROCESS OR ALLOCATIONS IS PROVIDED FOR ILLUSTRATIONS PURPOSES ONLY.

ANY STATEMENTS REGARDING CORRELATIONS OR MODES OR OTHER SIMILAR STATEMENTS CONSTITUTE ONLY SUBJECTIVE VIEWS, ARE BASED UPON EXPECTATIONS OR BELIEFS, SHOULD NOT BE RELIED ON, ARE SUBJECT TO CHANGE DUE TO A VARIETY OF FACTORS, INCLUDING FLUCTUATING MARKET CONDITIONS, AND INVOLVE INHERENT RISKS AND UNCERTAINTIES, BOTH GENERAL AND SPECIFIC, MANY OF WHICH CANNOT BE PREDICTED OR QUANTIFIED AND ARE BEYOND CFM'S CONTROL. FUTURE EVIDENCE AND ACTUAL RESULTS COULD DIFFER MATERIALLY FROM THOSE SET FORTH, CONTEMPLATED BY OR UNDERLYING THESE STATEMENTS

Other news

- Our Market Microstructure conference in Paris from the 6-9 December was well received
- We have submitted our paper Price impact without order book: A study of the OTC credit index market, to the Trading and Market Microstructure section of the Cornell University Library.
- Our paper Dissecting cross-impact on stock markets: An empirical analysis is to appear in Journal of Statistical Mechanic

Footnotes

[1] The trender used here is defined as the sign (either +1 or -1) of the difference of a 50 day exponentially weighted moving average (EWMA) and a 100 day EWMA

[2] Defined according to https://en.wikipedia.org/wiki/Relative_strength_index using 100 day exponentially weighted moving averages

[3] We use a CFM version of the Fama French implementation for momentum (UMD), value (HML) and size (SMB) and have tested the convergence with the data from Kenneth French's website. We note that other implementations, notably from brokers, are broadly in line with our conclusions for momentum and value. We hope to soon include discussion of a generic quality factor in this publication. Research is ongoing in this direction

[4] Generic 10 year yields as obtained from Bloomberg

[5] The TYVIX is calculated from the CBOT's options on 10 year futures, using the same methodology as the VIX, and is published by the exchange.

[6] The CBOE/NYMEX WTI Crude Volatility Index

[7] https://cib.societegenerale.com/en/sg-prime-services-indices/

[8] The EMN index is that calculated by HFR, while the CTA index is calculated by the Société Génerale

[9] For the EUR/USD exchange rate we use the Bloomberg defined EURUSDV1M ticker. The VIX index is calculated and published by the CBOE.

[10] We estimate effective FX volumes to be a factor of 5-10 more than this due to the extra liquidity available through the spot markets

[11] Liquid order book markets for FX exist but the non-order book volume dominates

[12] The Implementation Shortfall: Paper Vs Reality. Andre Perold, The Journal of Portfolio Management, 1988

[13] Both effects are accounted for in CFM's cost modelling but the details are beyond the scope of this text

[14] Costs (and the level of statistical uncertainty in cost) are often measured in basis points (bp) or hundredths of a percent, i.e. 1bp=0.01%. A cost of 10bp, for example, represents a cost of 0.1% of the face value of the share. For one share worth $100 this would be a 10 cent cost. A 1bp error is then a cost of (10±1) cents

[15] This average is similar in nature to a Volume Weighted Average Price (VWAP) except we are only including our own trades in the average

[16] This statement can be a source of confusion. We are talking about cost per share! The total dollar cost of a two share trade (2 times the cost per share) will be higher. It cannot be cheaper to trade more!

[17] It is also common to guarantee a benchmark fixing price such as one published by a central bank

Disclaimer

CFM has pioneered and applied an academic and scientific approach to financial markets, creating award winning strategies and a market leading investment management firm.

Contact us

Capital Fund Management S.A.
23, rue de l'Université,
75007 Paris, France
T +33 1 49 49 59 49
E cfm@cfm.fr

CFM International Inc.
The Chrysler Building, 405 Lexington Avenue - 55th Fl.,
New York, NY, 10174, U.S.A
T +1 646 957 8018
E cfm@cfm.fr

CFM Asia KK
9F Marunouchi Building, 2-4-1, Marunouchi, Chiyoda-Ku,
100-6309 Tokyo, Japan
T +81 3 5219 6180
E cfm@cfm.fr

Capital Fund Management LLP
64 St James's Street, London
SW1A 1NF, UK
T +44 207 659 9750
E cfm@cfm.fr