CFM

INSIGHT.DATA.CLARITY.

November 2016

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.

Contact details

Call us +33 1 49 49 59 49

Email us cfm@cfm.fr

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[1] 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, and therefore better price. These actions occur within milliseconds of each other on the most liquid markets and provide a rich source of trade data.

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.

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.

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

The cost of trading

From bid/offer spreads and commissions to impact

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.

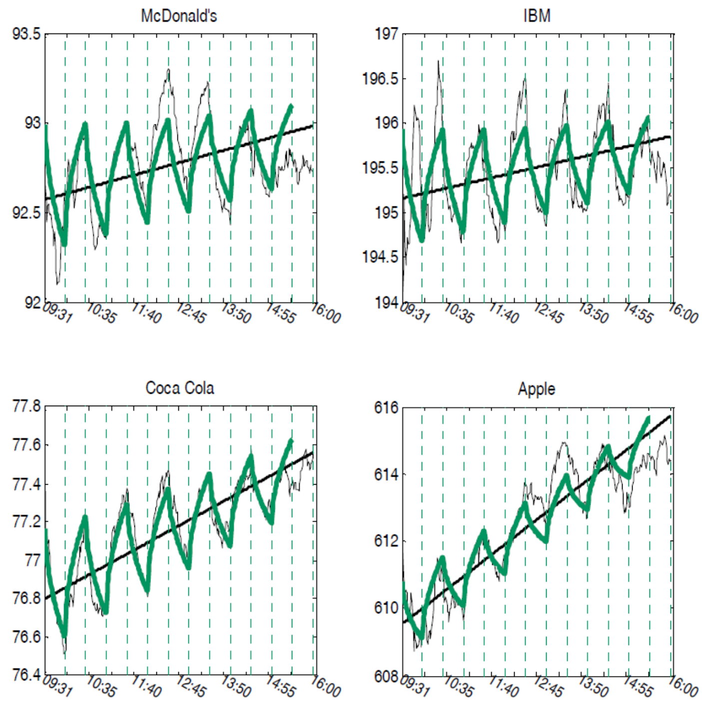

Figure 2: The price evolution of four stocks on the 19 July 2012. This oscillating, saw-tooth behaviour is more than likely caused by option hedgers in the market. Hedging options that are close to expiry and restricted to a small number of strikes (or even only one) can generate large sequential trades 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[2]. We assume that the meta-trade has been fully executed and that it is unpolluted from the impact of prior trades[3]. The implementation

[2] *The Implementation Shortfall:* Paper vs Reality. Andre Perold, The Journal of Portfolio Management, 1988

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

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[4] for typical stocks.

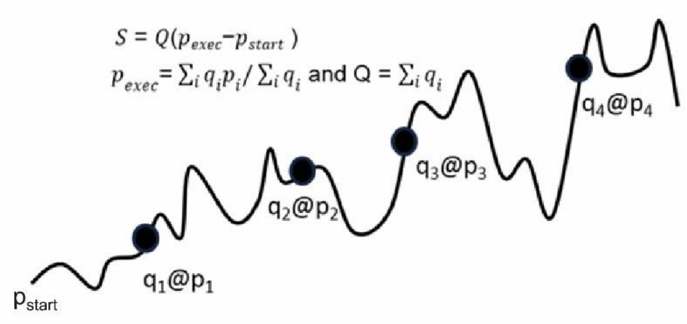

Figure 3: An illustration of implementation shortfall, the measure of costs used by CFM. A meta-trade of Q shares is split up into 4 trades of q1, q2, q3 and q4 executed at p1, p2, p3 and p4 respectively. The cost of each can be evaluated relative to the price p_{start}, before going on the market. The total cost of the trade is conveniently expressed as the difference between the quantity-weighted[5] 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 gets faster or 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

[4] Costs (and the level of statistical uncertainty in cost) are often measured in basis points (bp) of 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.

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

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 unit cost of a single trade[6]. 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 5 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.

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

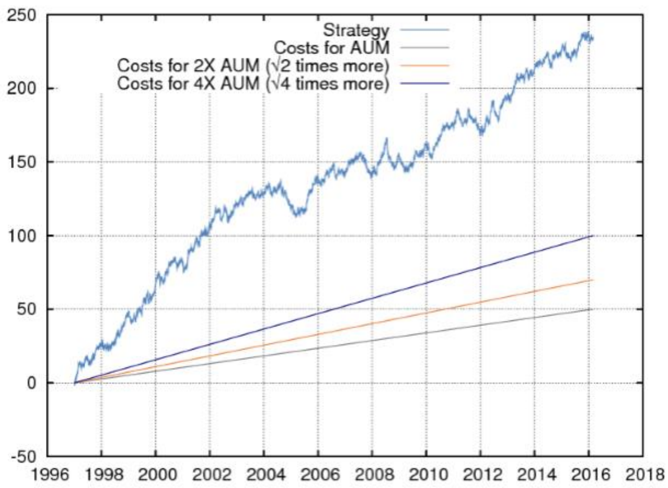

Figure 4: The "Profit and Loss" (P&L) curve for a fictitious trading strategy with a level of costs corresponding to a given level of Assets Under Management (AUM) allocated to the strategy. Also plotted is the increase in costs following a doubling of AUM, showing an increase in costs of √2=1.4, along 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[7] 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

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

that they have a guaranteed price but the broker's profit is actually increasing with volatility!

$$p_{exec} = \frac{\Sigma_i q_i p_i}{\Sigma_i q_i} = p_{close} - 1/3(p_{close} - p_{open})$$

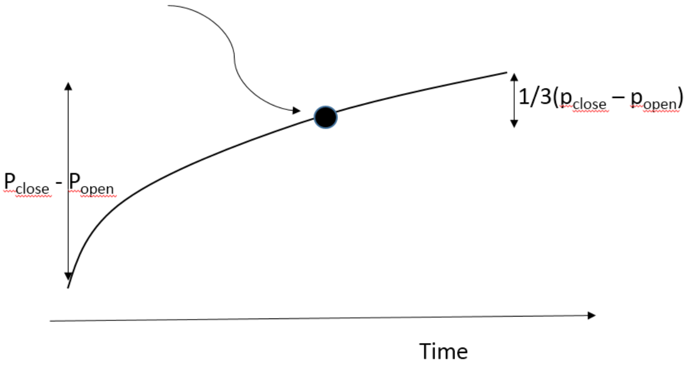

Figure 5: The average price evolution of a trade executed from the open to the close of a market day. The black dot represents the average execution price with the quantity acquired subsequently sold to the client at the close price. The broker pockets the difference between the two prices which represents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

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.

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