Technical Note

CFM
INSIGHT.DATA.CLARITY.

May 2018

IMPROVING TRADE EXECUTION SYSTEMATICALLY

Executive summary

The systematic execution of trading decisions by machines – as opposed to humans – can significantly improve the bottom line of an investment portfolio. However, computer code can be customised in an infinite number of ways, and in order to continuously improve such trade execution algorithms, one has to rely on data and statistics. This paper explains through concrete examples how following simple rules and monitoring the behaviour of different trading styles is instrumental in lowering trading costs and improving performance.

Contact details

Call us +33 1 49 49 59 49
Email us cfm@cfm.fr

Introduction

Investment management regularly involves decisions about one's portfolio. These are taken based on the prevailing market price at the time (we call this the decision price), and involve buying and selling different assets. These trades have to be executed on a listed or over-the-counter market, where the actual transactions will rarely happen at the exact decision price, because the market is moving permanently. They are done at a different execution price which is on average worse for the investor, i.e. higher for buys and lower for sells. In an earlier note we explained this effect in detail [1], here we will only say that in order to quantify such a difference, CFM uses the metric known as *implementation shortfall* (IS). This is defined as

(avg. execution price – decision price) x quantity,

with quantity taken as a negative number for sells. For example, if a decision was taken to sell 100 shares quoted at $30 on the market, but they were actually sold at a lower $29.90 on average, then

IS = ($29.90 - $30) x (-100) = $10,

this is how much the transaction has cost us.

Why does it cost money to trade?

Implementation shortfall has many origins [1]: brokerage and market fees, the bid-ask spread, and when doing several trades in a row also price impact (successive buys/sells pushing the price up/down). These numbers seem small, but when trading on a regular basis they add up quickly. If sufficient care is not taken to manage them, they can substantially degrade the performance of an investment portfolio.

The execution price is subject to market volatility, and so it intrinsically contains an element of luck. One can get a favourable market move on a trade, say if the price dips due to an announcement of poor earnings on the stock just before buying, so it can be obtained at a lower price. However, if one wanted to sell instead, the same would be bad news. How much one is exposed to such random events depends on how often one trades. After many trades the element of luck averages out, and the tendency that execution costs money becomes clear.

Figure 1 below illustrates what typical graphs of annualised cumulative implementation shortfall would look like for portfolios of different kinds. The 'High turnover' portfolio makes several small trades every day on a large number of products, and it pays execution costs in a steady, predictable fashion that is very close to the benchmark formula discussed in [1]. The 'CTA-like' portfolio resembles doing one or two large trades per month on a lower number of products. One can see more noise in these curves, and even on an annualised level costs can vary from the benchmark. Finally, the curve 'Alternative beta' corresponds to a portfolio doing even fewer, very sizable trades per year. Under these conditions chance plays a major role, and annualised implementation shortfall can occasionally even be negative, (execution actually making money by chance) but also easily double or triple what was expected.

How does one reduce trading costs?

There are several straightforward ways to reduce trading costs in a portfolio, including:

- Negotiating lower commissions with one's brokers
- Looking for cheaper products to trade: listed derivatives instead of OTC, futures instead of ETFs
- Concentrating trading on the most liquid products and the more active hours of the day
- Being more patient and allowing for slower, less aggressive execution.

Figure 1: Examples of annualised cumulative slippage for portfolios of different kinds.

However, eventually these options become exhausted, and one has to search for savings elsewhere. Algorithmic trading is an efficient way of further reducing costs, while also mitigating operational risk. Portfolio decisions are electronically sent to one or more computers that are themselves connected to the exchanges and platforms where trading takes place. These computers then take decisions about the exact sizing and timing of individual orders in order to execute the portfolio decision. The trading algorithms doing so can be customised in an almost infinite number of ways: timing trading with respect to market liquidity, using passive or aggressive orders, the choice of markets for sending each order, etc. Faced with so many options regarding trading style, can we learn at all which one to pick? Which precise algorithm will work best?

Ideas for designing execution strategies can be found in several places. There exists plenty of directly useful academic research on market structure and costs [2], [3]. High-resolution market data can also be analysed to gain additional insights. With the proper software and hardware infrastructure one can even replay pre-recorded real-time market data and run market simulations that resemble actual trading very closely, in order to evaluate performance in 'paper trading'. This is similar to back-testing procedures that are becoming increasingly commonplace today, except instead of using daily prices, calculations can be timed with a precision of a millisecond.

The most important thing is, however, to analyse one's own past trading and its performance. No amount of theoretical calculation or computer simulation can fully replace actual experience. In the following we will explain how one can turn data about past executions into a tool for evaluating different ideas. However, let us first digress to introduce the *multi-armed bandit*, a mathematical concept well-known in diverse domains ranging from web advertising to the organisation of clinical trials.

A detour to Las Vegas

Imagine you are a gambler and you decide to go to Las Vegas to play slot machines (which each have an arm, hence the name multi-armed bandit). When you enter the casino you are confronted with lots of similar looking machines. Some of them may provide better odds than others, but there is no way for you to know in advance, because other gamblers do not let you to watch them play. Thus you will only have your own experience, your own losses and gains to learn from.

In order to figure out which machine is 'good' and which is 'bad' you will have to play each of them multiple times, but there is still a large element of chance. It is difficult to distinguish a bad machine from just bad luck while playing a good machine. How do you learn both which machines are the best, and make the most money at the same time?

You would do well to keep in mind a few easy-to-follow rules.

Rule #1: Above all, it is important to set realistic expectations. Figure 2 shows hypothetical 'winnings' over several years of playing two different machines, which are possibly negative. This is expected; we know that slot machines are set up such that on the long run the guest loses money and the casino wins.

Nevertheless, we see a small difference between the two machines, compared to a 'benchmark' of losses which would be expected from an average machine, as shown in Figure 2:

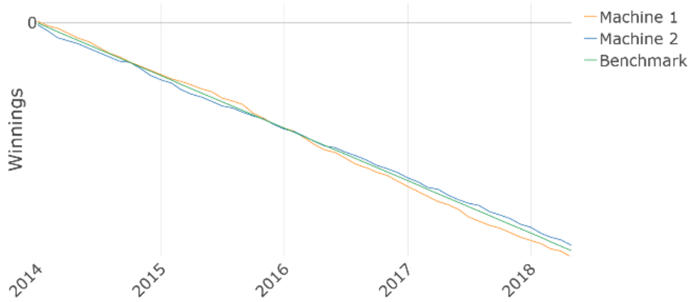

Figure 2: Two losing machines versus the benchmark.

In fact it is intuitive to look at performance with respect to such a benchmark. Not only does this make for a nicer graph, but it also eliminates the illusion that the two machines are the same because the lines go roughly together. Subtracting the benchmark for losses we see the difference in Figure 3:

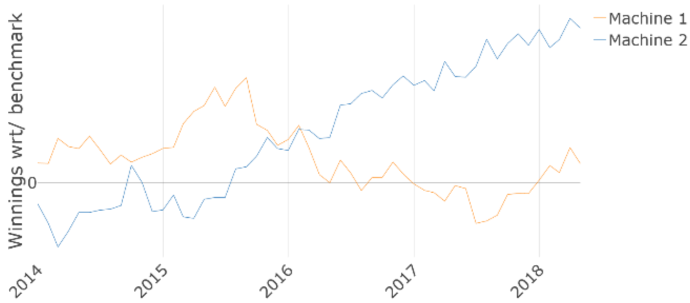

Figure 3: Performance after subtracting benchmark.

Rule #2: Let us take another example, again two machines that at first sight have the same final performance:

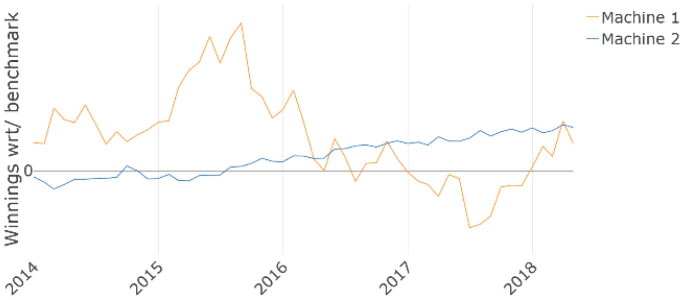

Figure 4: Performance of two machines with difference bet sized.

However on closer observation one notices that the orange line is much more jagged and its trend less consistent. This is because we have actually played Machine 1 with five times larger bets than Machine 2! The comparison is thus unfair, relatively to the amount we have bet the two do not perform the same at all! If we divide the first curve by five to compare fairly, the difference becomes apparent:

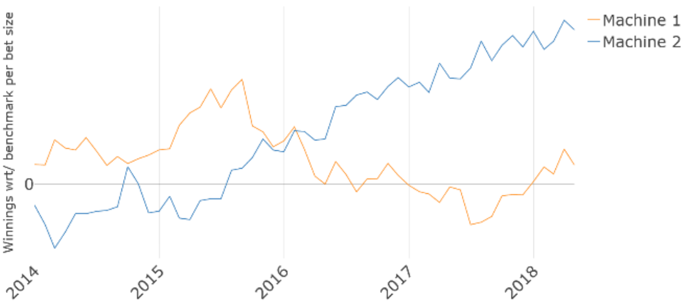

Figure 5: Performance after normalisation by bet size.

Rule #3: In real life statistical differences only reveal themselves over time. We show in Figure 6 the behaviour of four different machines:

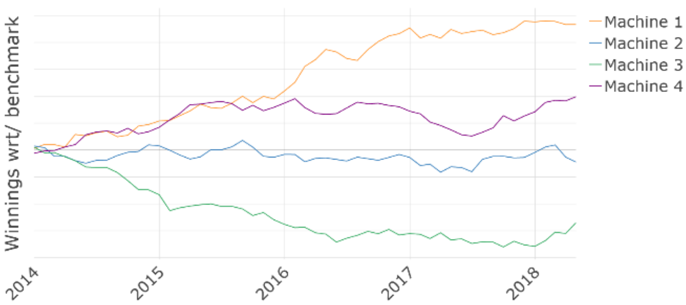

Figure 6: Performance of four machines compared.

In Figure 6 we can see that Machine 1 is best. However, we could not have known that at the end of 2014! If we only have data up to that point, Machines 1 and 4 look about the same, and 2 is not very far behind. By the end of 2015, number 3 looks clearly worse, but we still could not tell the difference between 1 and 4. One has to be patient and gather a sufficient amount of data so one is not fooled by statistical flukes in these comparisons.

How is all this relevant to trade execution? Let's have a look at that next!

A systematic approach to reducing trading costs

In systematic trading, styles of trade execution are formalised as computer code. These are often called 'execution strategies'. Since a priori we cannot know with certainty which style will work best, i.e. generate the lowest execution costs, we must rely on experience. As discussed at the start of this paper, just like playing in a casino, having to execute trades generates losses on average. Realistically our aim must be to – over time – continuously improve our strategies in order to pay less and less. Regardless of the change of context, the same rules should apply as for slot machines.

Rule #1: Use benchmarks for your expectations. Look at how much executing via a given strategy saves with respect to expected cost, for example the formula explained in [1].

Rule #2: Compare costs relative to trading activity. If you ask a strategy to trade twice as much, you should allow it to spend twice as much in costs.

Rule #3: You need time to learn. The performance of the different strategies has to be followed over time, and conclusions should only be drawn based on sufficient data. This helps to avoid making decisions based on just chance.

Naturally the more often one trades, the quicker one gathers data. As illustrated in the first section of this paper, this means that for high-turnover portfolios actual execution costs conform well to the expectations as quickly as a few months. When only executing low-turnover, for example alternative or smart beta type portfolios which trade infrequently, one may have to wait several years to reliably measure execution quality.

A concrete, real life example

In the above example we advocate an active, experimental approach for comparing trading styles, it is time to put these ideas to the test. To illustrate this, we will show the actual cost savings generated by CFM's execution strategies competing on the UK equity market, in the series of figures below. This is a small, but typical segment of all trading activity.

The savings were adjusted according to the rules explained above to make them comparable: they are with respect to the benchmark and relative to trading volume. Hence the curve going up fastest will correspond to the best strategy. Let us now use this data to tell our story of continuous improvement. We will gradually reveal performance data as collected by CFM, in each step followed by an explanation.

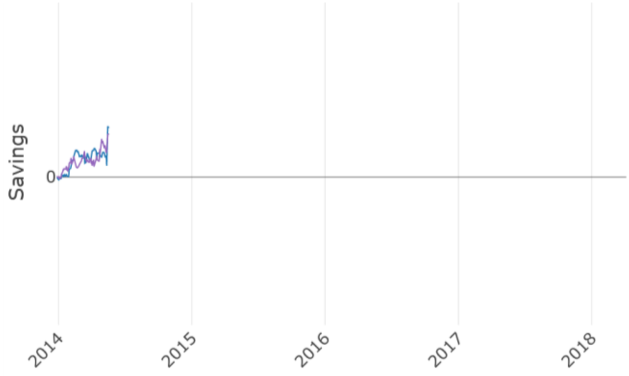

Figure 7: Comparison of strategy performance mid-2014.

In early 2014 the UK market was executed using two competing strategies, for simplicity let us call them 'blue' and 'purple'. Their performance was similar: the two savings curves increase at the same speed.

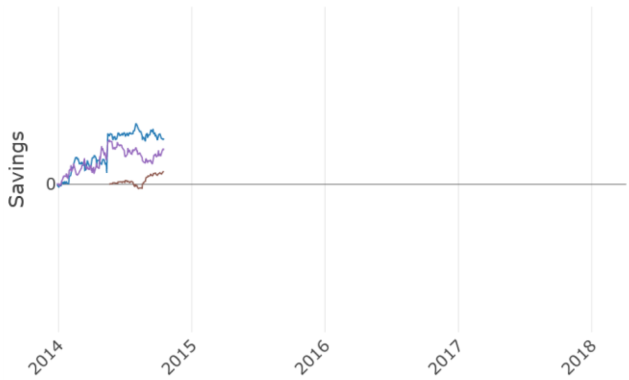

Figure 8: Comparison of execution strategy performance at the end of 2014.

In mid-2014 we introduced a new improved 'brown' strategy as seen in Figure 8, but from the initial few months of data it was not clear whether it performed better than the previous strategies, which we have also kept active.

Figure 9: Comparison of execution strategy performance late 2015.

In parallel, based on some new ideas we introduced the 'green' strategy as shown in Figure 9 at the end of 2014. We allowed all four to run in competition for some time, while working more actively on other markets. Over 2015 the 'green' strategy had disappointing performance, and it was stopped. Abandoning a strategy is not decided upon lightly, and we relied on additional results in research that showed *why* the strategy fared poorly, rather than relying only on the pure statistics that it does.

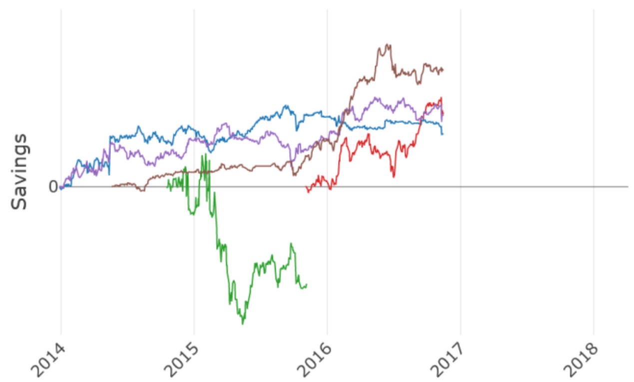

Figure 10: Comparison of execution strategy performance late 2016.

This research also allowed us to replace the 'green' strategy by a better 'red' one as shown in Figure 10. At the same time the 'brown' strategy showed signs of outperforming the original 'blue' and 'purple'. This made us grow confident enough to start phasing out old strategies, starting by decommissioning the oldest 'purple' one.

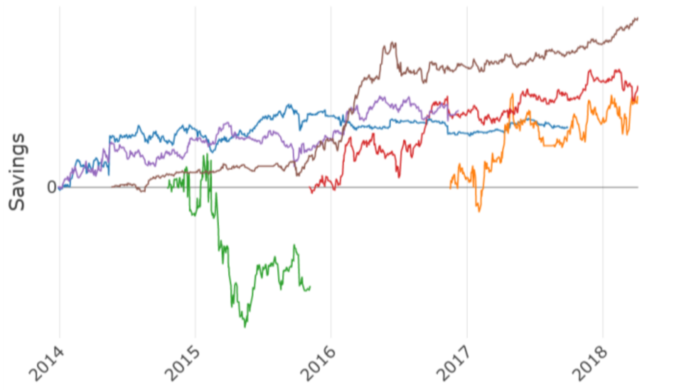

Figure 11: Final comparison of execution strategy performance early 2018.

The 'red' strategy continued to outperform through 2017, and based on promising research results we gradually rolled out a new 'orange' one, as seen in Figure 11, which performed even better through 2017. However, one can see that the strategy did have an early dip in performance due to poor luck, and it was important let it run further to gather statistics. The remaining other old 'blue' strategy was also finally decommissioned.

Over several years we have thus confirmed three ('brown', 'red' and 'orange') useful improvements to our execution style.

These and even newer strategies continue to compete. Along with results obtained on other products and markets, they can also be combined in future algorithms to obtain even higher cost savings.

Take-home message

Active research and development in algorithmic execution has significant upsides. Lowering execution costs and better management of the capacity of one's portfolios is an ambitious, but very important goal. In the above we have presented the basics of our empirical approach to achieve this goal.

This work requires a large amount of trading, and the meticulous gathering of data. Moreover, because this process takes months at best, and sometimes years, we prefer to complement it with continuous research, and above all common sense. We firmly believe that by collaborating with portfolio managers and continuously improving execution style one can, at the end of the day, deliver better performance to the final investor.

References

[1] Capital Fund Management, 'Executing With Impact: Why the price you want is not the price you get!' 2016.

[2] Jean-Philippe Bouchaud, Julius Bonart, Jonathan Donier and Martin Gould, Trades, Quotes and Prices: Financial Markets Under the Microscope, Cambridge University Press, 2018.

[3] Charles-Albert Lehalle and Sophie Laruelle, Market Microstructure in Practice: 2nd Edition, World Scientific Publishing, 2017.

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