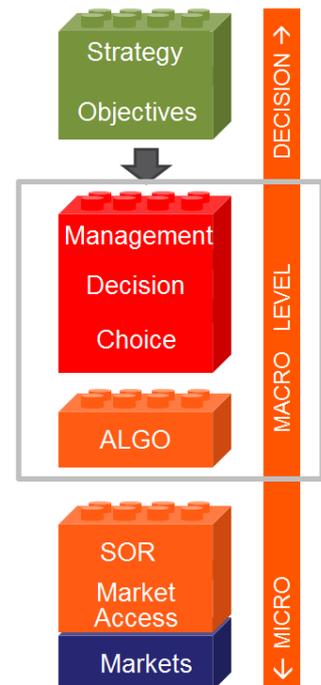


Execution overview – The macro level

Compared to the micro level that we described earlier in “The micro level” document, macro is, as the word indicate execution on a higher level – and it is closer to where it all starts; the actual strategic **decision** to buy/sell.

Depending of the size of the investment and the available liquidity – the execution objectives have to be selected. These objectives serve as input at the **macro level** of execution and typically involve dividing a large investment decision into smaller, more manageable parts. Each smaller part is then transmitted to the micro level of execution, which involves access to markets and Smart Order Routing.

What we refer to as the **macro level** of execution is the way we choose to execute in and out of investment positions – and is the focus of this document.



The cost of trading

There are explicit costs of trading, such as crossing spreads, commissions, taxes and fees. These are all easy to benchmark. There are also implicit costs in the form of market impacts and market movements that are trickier and demand special attention.

Market impact

Market impact is the effect (on the market) that a market participant has when he/she buys or sells an asset. It is the extent to which the buying or selling moves the price away from the buyer or seller. A financial institution that is seeking to manage its market impact needs to limit the pace of its activity.

Market movement (market risk)

Market movement is the risk (in time) that the market will move in an unfavorable direction, against your intentions. This risk increases with the time spent executing an order.

Decision – the tradeoff between market impact and market risk

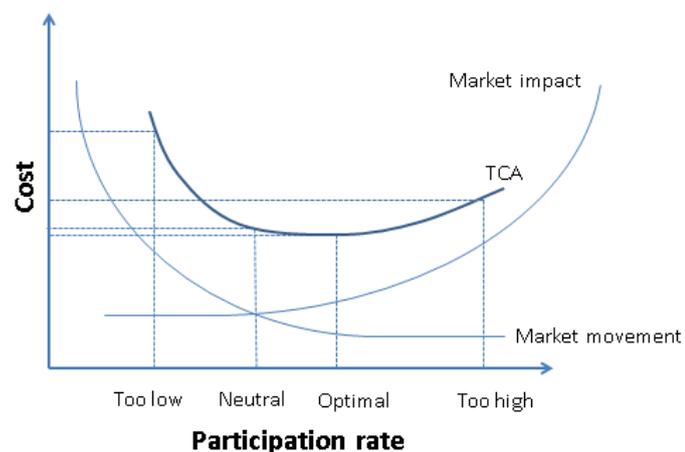
One major difference with macro execution compared to micro execution is that the time horizon is significantly longer (up to several days). During this time, there are many potential event/actions/factors in addition to our own actions that could occur in the market and affect our execution. If we, pre-trade, have an opinion or feeling about where the market is heading,

we may choose to trade with a short time horizon and can thereby cause market impact. If we believe that the market will move in a favorable direction we may simply wait and not trade at all. Most of us don't have access to this short-term market information and are simply seeking to trade in and out of position in a neutral and balanced manner. Avoiding market impact becomes a key factor. Studies show that the optimal participation rate is somewhere between 2% and 15% (depending on factors like instrument liquidity, size of order and volatility).

Executing with a too high participation rate may lead to market impact. A too low rate takes an unnecessarily long time, causing market movement risk (i.e. loss of alpha). See picture at right.

The tradeoff between market impact and market risk (time) is an important decision point for managing the investment.

This decision is the input for the macro level of execution.



Benchmarking macro execution

From an execution service provider perspective, one important factor in benchmarking is the arrival price, which equals the price of the instrument when the order arrives to us. Read more about best price, European Best Bid and Offer in our micro level document.

Common benchmarks used:

- Implementation shortfall (IS) – the difference between the arrival price and the final execution price (price including commissions, taxes, etc.) for a trade. Also known as “slippage”.
- Market impact – How much did our execution move the price of the asset?
- Spread capture – The ratio of passive and aggressive fills.
- VWAP – The Volume Weighted Average Price during the time of execution. Often called “market average”.

Problems with different benchmarks

- Implementation shortfall – The drawback is that it includes the effects of all other participants' actions and all other significant events occurring in the market during our execution (noise). It is difficult to separate our own effect (impact) on the market from this noise, most would say impossible on a single order. Adjustments using indices or sector indices may be used in an attempt to remove noise.

- Market impact – Market impact is difficult to benchmark and there is no de facto standard, although there are some actions that clearly have a market impact and should be avoided if not part of a well-informed and planned decision. Examples of such actions are:
 - Executing at a too high participation rate, being an excessively large part of the market.
 - Placing a relatively large order that is totally visible in the order book
 - Constantly reposting a new order of the same size if one order is filled – a pattern that is easy to detect and sends signals.
 - Buying several levels at one time, especially if the levels do not return (reload).
 - Entering relatively large orders in a dark pool, with a minimum accepted quantity that is too small (allowing pings – others finding the order using small test orders).

- VWAP – The final execution price is compared to the market volume weighted average price for the duration of the order. This is easy to explain and to implement and is commonly used as a benchmark for the “neutral execution style”. The drawback is that it is relatively easy to game, meaning that market impact should be monitored at the same time.

- Pre-trade estimate – Through theoretical models it is possible to calculate an estimate of the final execution price, which is then compared to the execution price. This is associated with the same drawback as implementation shortfall, since there are many unpredictable factors that can affect the result.

- Spread capture – Passive execution captures one spread. Aggressive execution crosses the spread – with no capture. A high spread capture indicates a good venue selection when placing passive orders in the market, and that we are not aggressively pushing the price. Excessively large passive orders in the market may capture the spread, but might also signal your intentions (passive volume has more pre-trade visibility than aggressive). Again, market impact should be monitored at the same time.

Summary of benchmarking problems

Some benchmarks such as VWAP and spread capture are more focused on measuring the execution when we have no pre-trade view of where the market is heading. We want to execute in line or better than the market, “beat the market average”. For other benchmarks, such as implementation shortfall, the trick is to separate our execution from the market noise. If the market goes up and we are buying, our implementation shortfall will be bad, no matter what. If it is falling, we will look good (if we are executing with an algorithm like VWAP). One solution is to analyze huge numbers of orders over a long period of time using advanced cause and effect models.

Algorithms

Generally, all algorithms have built-in anti-pattern recognition functionality, for example randomizing order size and the time between orders to avoid sending signals that might cause market impact. The choice of algorithm is closely linked to the choice of benchmark, and some algorithms even have the same name as the benchmark in question. TWAP, VWAP and Participate are all very similar, they are just different ways of entering and following a participation rate that is either assumed to be linear, following a historical pattern, or trying to follow current market conditions.

TWAP (Time Weighted Average Price) – Splits a large order into smaller equal-sized parts or slices that are spread over the selected time duration. TWAP doesn't use any statistics and will not adjust for any events during the day (linear).

VWAP (Volume Weighted Average Price) – Follows the market's normal volume distribution based on historical data.

Participate – Executes at a set participation rate, for example 15% of the market. In other words, if someone trades 100k – within a reasonable amount of time and with some variation and lag to avoid signaling – I should trade 17.5k (includes our own volume).

IS (implementation shortfall) – In simple terms, this removes the balancing of market impact and market risk from the customer/broker and leaves the choice to the IS algorithm. Predicting the movement of the traded instrument becomes a key issue, since if we are buying and the price is falling, we should most likely wait and buy later. If rising, we would rather buy all now and be done/complete.

Dark – By executing an order using so-called dark pools, the user can enjoy a minimum of pre-trade information leakage and minimized market impact. Special attention must be given to setting a Minimum Accepted Quantity (MAQ) in order to avoid gaming (so-called pinging). Dark is commonly used by other algorithms which have volumes that are not active in the market. If we can execute/complete the order at a specified price, we reduce the market risk of trading.

Sample of execution and benchmarking using LiquidMetrix

Trading Summary

To the right we can see that during execution of this VWAP, the Average On-book Spread (EBBO spread) was 6.92 bps (1 bps/basis point = 0.01%). We traded for 1h and 36m. The participation rate of 1.59% is low; execution was 77.9% passive, and 3.7% dark. Overall, this indicates a low market impact. Execution was carried out using MTFs to a high degree, indicating a lower cost of trading compared to the primary market (this should be weighed together with execution quality).

Trading Summary

Value Traded:	€134,306		
Volume Traded:	27		
Number of Fills:	19		
Average Fill Size:	€7,068.71		
Duration:	1h 36m		
ParticipationRate:	1.59%		
Book Spread (depth):	6.49 BPS		
Book Spread (touch):	6.92 BPS		
Execution Style:	Aggr	Pass	Mid
	18.4%	77.9%	3.7%
Venue Types:	Prim	MTF	Other
	0.0%	96.3%	3.7%
Venues Traded:	CHIX: 85.10% BURG: 11.17% CHID: 3.73%		

Order Execution History

We are selling but the stock price is falling fast (the market is moving away from us). Our volume distribution in % is represented by the green line, and is reasonably stable.



Performance Summary

Spread capture (SP) is generally considered a good benchmark for successful algorithmic trading. This algorithm had an SP of 98.03%. The implementation shortfall is -63.928bps, which might seem terrible. But we are selling in a falling market. If we had had pre-trade knowledge about a falling market, we would most likely have sold using a high participation rate instead of a VWAP (thereby most likely adding to the speed of the falling market). The VWAP benchmark is 15.934bps, which means that we have beaten the market average by almost 16bps. It becomes clear that choice of strategy and benchmark is more important than anything else.

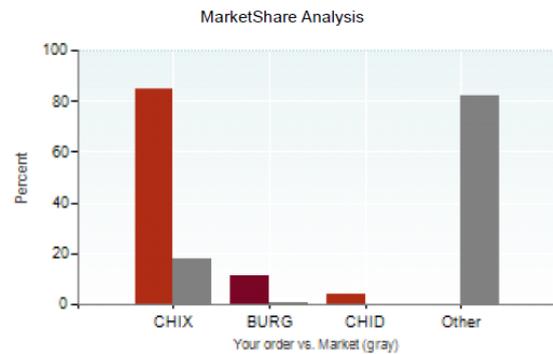
Performance Summary

Imp. Shortfall (First Fill):	-63.928 BPS
VWAP (First To Last):	15.934 BPS
Avg. Spread Capture:	98.03%

Benchmark	Price	(+/-) BPS
Price Achieved	36,963.700000	
Arrival Price (touch)	37,200.000000	-63.928
Arrival Price (mid)	37,220.000000	-69.338
First Fill Price (touch)	37,200.000000	-63.928
First Fill Price (mid)	37,220.000000	-69.338
Previous Close	37,600.000000	-169.229
Opening Price	37,360.000000	-106.076
Closing Price	36,800.000000	+44.484
VWAP (First To Last)	36,904.894366	+15.934
VWAP (Daily)	37,115.667177	-40.944
VWAP (20% Part.)	37,261.037037	-79.798

Market Share Analysis

The graph at right shows the venue distribution of the execution. This is actually a micro level and not a macro level part of execution. However, the high spread capture of 98.03% indicates that our passive posting has been successful even if we used only MTFs in this case and did not make a single trade on the primary market. Since our execution is well benchmarked, we know that the cost savings from using MTF trades had no negative effect on execution quality. For some instruments (like this one), quality increases when trades are executed on alternative markets.



Venue Breakdown

A deeper look into where we executed and how.

Success (%) is important, this indicates success at trading within EBO (see micro level document).

Venue Breakdown

Venue	Trades (#)	Order (%)	Market (%)	+/-	Success (%)	SOR (BPS)	Imp (BPS)	SC (%)
CHIX	15	85.10	17.76	-67.34	100.00	0.00	4.47	75.4
BURG	3	11.17	0.21	-10.97	100.00	0.00	14.35	160.0
CHID	1	3.73	-	-3.73	100.00	0.00	5.37	50.0

Summary

TWAP, VWAP and Participate are basically different ways of participating using linear, historical or actual market volume as a base. They all fit neutral execution and execution strategies where we have a pre-trade feeling of where the market is headed. Only parameters and benchmarks will vary with execution style. The IS algorithm is an automated way of trying to analyze the market and making the balance of market impact and market risk automatically beat the neutral execution style. The problem here is to prove the value, taking into account all of the noise in the market. Everything done at the macro level is dependent on how well the micro level performs the actual trading. Both layers have to be well tuned, tightly integrated and benchmarked to achieve a good overall result in line with your intentions. All machine-handled algorithms add some kind of rules and pattern to the trading that can potentially be detected (there are patterns even in chaos). Making frequent changes in settings is one part of the trimming process, but is also part of an anti-gaming strategy.

At Neonet, we strive to deliver a truly transparent and independent execution service with an optimized balance of quality and cost. Transparent reports like the ones illustrated in this document give our clients full control and insight into how we perform.

For more information, visit www.neonet.com.