MARCH 2019

ANALYSIS OF THE AGGRESSIVE BEHAVIOUR OF MARKET PARTICIPANTS: DO HFTS TRADE OPPORTUNELY?

PAMELA SALIBA

Risks & trends

AUTORITÉ DES MARCHÉS FINANCIERS

AMF

This report was prepared by the Market Surveillance Division. It is based on sources that are considered to be reliable, but whose comprehensiveness and accuracy cannot be guaranteed. The views expressed in "Risks and Trends " are those of the authors; they do not necessarily reflect the position of the AMF.

This report may not be copied, distributed or reproduced, in full or in part, without the express, prior, written authorisation of the AMF.

ANALYSIS OF THE AGGRESSIVE BEHAVIOUR OF MARKET PARTICIPANTS: DO HFTS TRADE OPPORTUNELY?

This note aims at describing the aggressive trading behaviour of market participants, and especially High-Frequency Traders (HFTs). It is commonly accepted that HFTs often carry out market making activity, and therefore submit mainly passive orders. However, it seems important to point out that, even when they carry out market making strategies, these are not purely passive strategies and they also send aggressive orders for inventory management and in the course of their overall trading strategies.

We study price evolution around aggressive orders to identify market participants who trade the most opportunely. With this approach, we can measure the informational advantage of aggressive trades, which are usually considered as more informed than passive ones. This study therefore assesses the validity of the hypothesis that HFTs are better informed than the rest of the market, and hence are capable of exploiting arbitrage opportunities, and generating high profits whilst minimising the level of risk taken.

We find out that HFTs market makers belong to the 25% market participants realising the highest short term potential profits. This likely shows that they succeed in realising a smart inventory management and using efficiently aggressive orders in their trading strategies.

Moreover, for a given same market participant this analysis allows us to identify and compare the various strategies followed by a same member reference code, such as high-frequency trading strategies seeking short term potential profits and longer term strategies. This measure of potential profit/gain has several advantages: not only does it provide a means of quantifying informational advantage, but additionally it provides a complementary measure for identifying HFTs, and is relevant to pinpoint the different strategies used by traders, such as mean reversion, trend following or another particular strategy consisting in seizing opportunities, such as orders inserted inside the bid-ask spread, faster than other market participants.

1. DEFINITION

1.1. TRADING OPPORTUNELY: WHAT DOES THIS MEAN?

Every transaction taking place in a continuous Limit Order Book (LOB) is the outcome of the matching of an aggressive and a passive order. An aggressive order triggers the transaction: this is the most recent order submitted in the LOB[1] before the transaction, whereas a passive order is the order that awaits execution in the LOB.

Trading opportunely involves the sending of aggressive orders to the LOB at the right moment and in the right direction with the object of realising a profit.

In order to asses a trade's opportuneness we therefore analyse price evolution over multiple time horizons before and after each aggressive order and compute a potential profit based on the trade price and the price change. If the price goes up after a buy trade, then the potential profit is positive and the trade is opportune, otherwise the potential profit is negative and the trade could have been executed more opportunely at a different moment in time.

As HFTs use more sophisticated infrastructure and automation technologies than other participants, it is therefore quite legitimate to expect them to trade the most opportunely. We aim at evaluating potential profits from their aggressive trades and compare them to those from the rest of the market. This enables us to identify market participants who trade both aggressively and most opportunely[2], according to the different time horizons. Note that potential profits are presented without taking platform fees and settlement costs[3] into consideration.

1.2. WHAT IS HIGH-FREQUENCY TRADING (HFT)?

The HFT is a subset of algorithmic trading[4] for which minimising latency is a crucial element for performance. HFTs may use co-location and proximity services to minimise latency. Most of them submit large numbers of orders that are cancelled relatively shortly after submission, trade large volumes, consistently maintain a low inventory level by holding positions for very short time and turning them over very rapidly.

In order to identify HFTs, the AMF relies on its knowledge of market participants and on various metrics, including a metric based on the lifetime of cancelled orders. This classification uses two sets of conditions:

---

[1] As an exception, a stop order which has been triggered and which has turned into a market order may have been entered before the passive order it trades against. However, it can also be considered in this particular case that the market order has been submitted by the order management facility of the exchange.

[2] The analysis of the opportuneness of aggressive orders could be completed, in a next step, by an analysis of the passive orders in order to identify participants who undergo the lowest adverse selection. For the time being, we focus on aggressive orders only. See appendix for further details about the adverse selection.

[3] Taking fees and settlement prices into consideration can impact the net profit realised by each member, especially since not all members are charged the same level of fees by the exchange, and only members holding end-of-day positions incur settlement costs. For instance, market makers have advantageous fees and do not pay settlement costs, since they have a net flat position at the end of the day.

[4] MIFID II states that algorithmic trading means trading in financial instruments where a computer algorithm automatically determines individual parameters of orders such as whether to initiate the order, the timing, price or quantity of the order or how to manage the order after its submission, with limited or no human intervention.

- Condition 1 is based on a comparison with other participants: the participant must have cancelled at least 100,000 orders during the year, and the average lifetime of its cancelled orders should be less than the average lifetime of all cancelled orders in the LOB
- Condition 2 is based on a set threshold: the participant must have cancelled at least 500 000 orders with a lifetime of less than 0.1 second (i.e. the participant quickly updates the orders in the limit order book) and the top percentile of its cancelled orders must be less than 500 microseconds (i.e. the participant regularly uses high-speed access to the market).

Note that HFTs can follow different strategies not limited to market making, such as other mean reverting strategies or trend following. In the following part of the study, the strategy of HFT market makers will be studied specifically and compared to that of other HFTs.

2. DATA & MEASURE OF THE OPPORTUNENESS OF AGGRESSIVE ORDERS

2.1. DATA DESCRIPTION

The analysis is conducted on the CAC 40 stocks traded on Euronext Paris over a 3-month period, from September 2017 until November 2017, during which the volatility on the CAC 40 was stable and reached historically low levels[5]. Furthermore, this studied period is neither disrupted by end of year trading effects nor by MIFID II that entered into force in January 2018. Note that this study focuses solely on the analysis of strategies on Euronext, and does not consider other trading venues. The whole data set contains approximately 8 million aggressive orders and 423 million events[6].

The study focuses solely on aggressive orders: one single aggressive order triggers one or several transactions. Over the period analysed, we used both trade data and LOB data to describe the evolution of the LOB precisely before and after each aggressive order for each market participant. How the best bid and ask prices build up and evolve in the LOB allows us to assess the potential profit of each aggressive order. Note that we remove market data corresponding to the initial and final twenty minutes of the trading session, as these periods usually have specific features because of the opening/closing auction phases.

Note that Euronext provides the AMF with information concerning each participant involved in a transaction or order submission, in particular:

- Member code
- Trading capacity: liquidity providing including SLP (Supplemental Liquidity Provider) and RLP (Retail Liquidity Provider) members, proprietary trading, agency trading or RMO (Retail Member Organization) member. SLPs are market makers committed to meet Euronext's specific requirements of liquidity providing. The RMO (Retail Member Organization) and the RLP (Retail Liquidity Provider) exist as a part of a programme offered by Euronext: the Retail Matching Facility specialised in providing liquidity for retails. As RLPs can trade only against RMO orders stemming from retail participants and are therefore protected against adverse selection risk, they can offer tighter prices. As a matter of fact, it turns out that all RLPs are HFTs market makers.
- The SLE (Serveur Local d'Emission), which is a connection linking the member to the exchange infrastructure[7].

[5] See in the Appendix the comparison between the implied volatility during the analysed period and the preceding period (since 2013).
[6] An event can be either an order insertion, an order cancellation, an order modification or a transaction.
[7] An SLE can be used by different member codes (all corresponding to the same member participant), and for different types of activity. This service is fee-paying; its price depends on its capacity, which is defined in terms of the number of messages that can be supported per second.

Potential profit of aggressive orders can therefore be estimated at each level of granularity[8]. Since a same participant can have several member codes, and can hold various trading capacities through different SLEs, potential profit will be estimated accordingly:

- First we incorporate the trading capacity in order to establish if there are any disparities relative to the same member code according to its various trading capacities.
- With the awareness that some participants dedicate certain strategies to a particular SLE, we take the SLE into consideration in order to establish if there are any disparities relative to the same member code holding the same trading capacity, according to the various used SLEs.

2.2. THE MEASURE

This study aims at quantifying the opportuneness of an aggressive order by measuring its potential profit. This profit is estimated on the basis of the evolution of the gross price range around each aggressive order. For a given buy aggressive order, the profit is estimated as the difference between the price at the best ask at a given point in time (considered as the price at this moment at which at which the position can be unwound passively), and the volume weighted average aggressive order price[9]. By way of illustration, if the price executed by the buy aggressive order is equal to 10 euros, and at a given moment (after or before the aggressive order), the price at the best ask is equal to 10.1 euros, the difference is then equal to 0.1 euros. We note that, in this note, explanations are illustrated using buy aggressive orders examples, but in practice, buy and sell aggressive orders are analysed simultaneously (for sell aggressive orders, the impact is multiplied by -1 because their impact on the price is quite symmetric[10]).

We can express the potential profit according to different types of unit: euros, ticks, percentage or spreads per stock[11].
We choose the spread because it provides a basis for comparison among all shares[12].

This measure is averaged over all aggressive orders, weighted by the executed aggressive orders volumes and computed according to a logarithmic time scale: from 17 minutes before until 17 minutes after the aggressive order). In fact, with the knowledge that HFTs realise their profits over short term horizons, we are *a priori* interested in analysing the potential profit of the market participant a few seconds after the aggressive order. However, we chose to extend the study to several minutes following the aggressive orders in order to study the persistence of this profit as a function of time. At a given point in time before (respectively after) the aggressive order, the potential profit indicates whether it was more or less advantageous to trade earlier (resp. later). Furthermore, at a given point in time after the aggressive order, the potential profit indicates the gain that the trader could make (or the loss that could be incurred) if the position were to be unwound passively at that moment on Euronext.

---

[8] The number of flows obtained depends on the level of granularity taken into consideration. Grouped by:
- Member name, it is equal to 108 flows.
- Member name and member code, it is equal to 117 flows.
- Member name, member code and activity type, it is equal to 190 flows.
- Member name, member code, activity type and SLE, it is equal to 496 flows. The number of SLEs varies between 1 and 31 by member.

[9] The average aggressive price is the average of all the prices triggered by the aggressive order and weighted by the traded amounts.

[10] See appendix for further details about the symmetry between buy and sell aggressive orders.

[11] The average spread is computed for each share, it is averaged among all the events happening between 9:20 and 17:10, and weighted by time.

[12] See appendix for further details about the sensitivity of the potential profit units to the tick size.

The market participants to trade the most opportunely are considered to be those who have the highest potential profits[13]. They can vary according to the time horizon under consideration.

Note: Autocorrelated aggressive orders sent by the same member can have an impact on the measure. For instance, the profit of a member can be artificially generated by autocorrelated aggressive orders sent by this same member: one member sending successively aggressive orders will temporarily impact the price, neither because of a trend following strategy nor because of a herding strategy. In the remainder of this study, we show that SLPs realise the highest potential profits. The autocorrelation results presented in the annex show that aggressive orders of SLPs are the less autocorrelated[14]. Therefore, we can conclude the high potential profits of SLPs are not artificially generated by autocorrelated aggressive orders.

3. PROFIT ANALYSIS

3.1. THE DIFFERENT TYPES OF PRICE IMPACT

The profit measure as described in the previous section depends on the size of the aggressive order compared to the amount available in the order book. Three categories need to be distinguished and are referred to as:

- **Partial aggressive orders:** they consume less than the quantity at the best limit hit by the aggressive order and do not have a direct mechanical impact on the price, but only modify the imbalance.
- **Exact aggressive orders:** they consume exactly the quantity at the best limit and have mechanical impact since they trigger a best price change right after the trade[15].
- **N-limit aggressive orders:** they consume more than the quantity at the best limit and have also mechanical impact since they trigger a best price change.

The following graphic shows the different states of the LOB until the first price variation, by distinguishing between the partial (case 1) and exact (case 2) aggressive orders[16]:

[13] We state precisely from the outset that we seek to identify the most significant potential profits instead of simply identifying positive potential profits. This is because, contrary to passive orders that undergo generally adverse selection, aggressive orders are more informed: due to a transaction, the price tends to increase (resp. decrease) if it is a buy (resp. sell) aggressive order.

[14] See appendix for further details about the autocorrelations relative to each trading capacity in the market.

[15] Within the exact aggressive orders, two different categories stand out. Consider for instance the buy aggressive orders. The first category (for instance, fill or kill orders) is constituted by aggressive orders whose size is equal to that of the best ask. Following this aggressive order, the best ask is no longer valid. The second category (for instance market to limit orders) is constituted by aggressive orders whose size is bigger than that of the best ask but at a price equal to that of the best ask. Following this aggressive order, the best ask becomes now the best bid. All along this study, we do not differentiate between these two categories because they impact the measure in the same way.

[16] We note that in our study, iceberg orders are not excluded. This does not impact the results, considering that aggressive orders consuming iceberg orders are extremely rare. For instance, on September 4th 2017, no aggressive order used iceberg features.

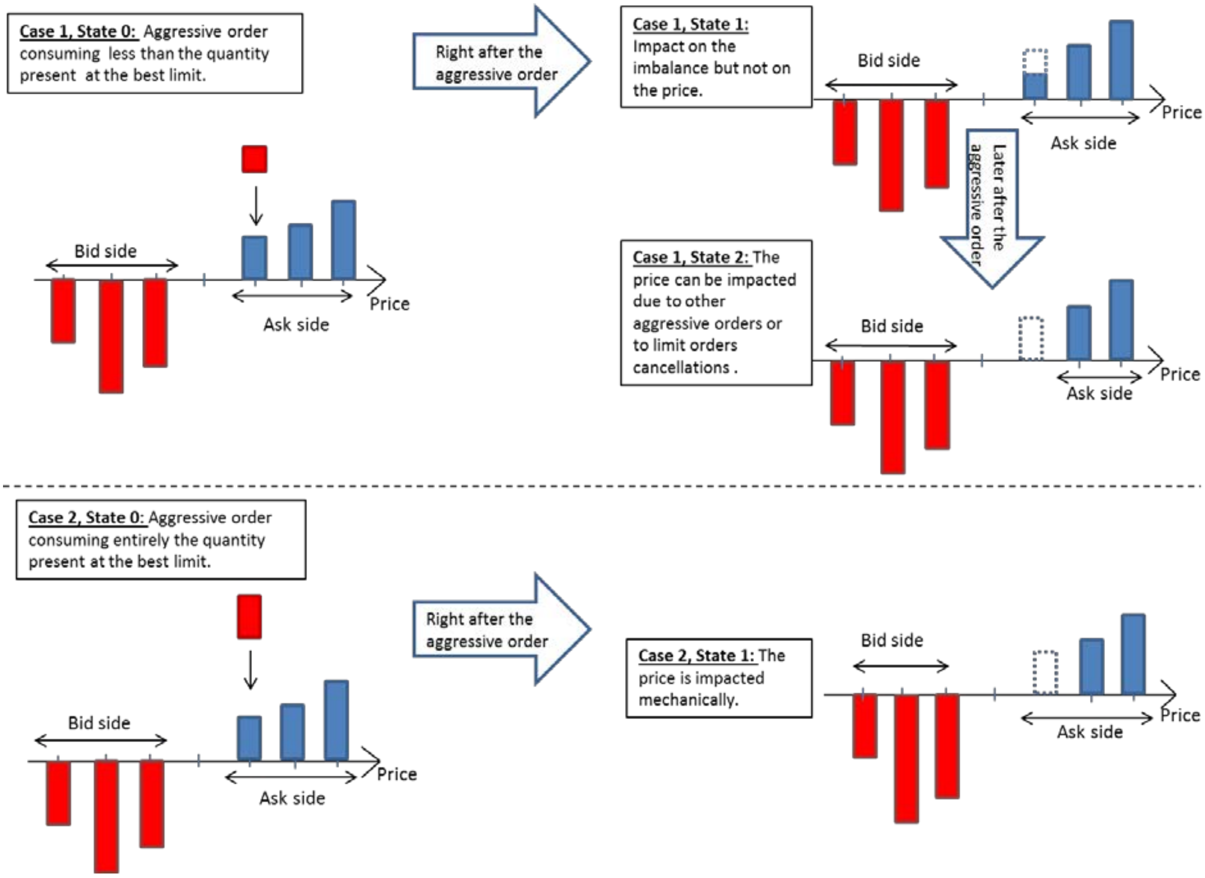

3.2. PRELIMINARY STATISTICS

The data set contains 8 million aggressive orders. The following table presents general statistics on the different aggressive order natures:

	Average traded amount per aggressive order	**Median traded amount per aggressive order**	**Share of aggressive orders number**	**Share of traded amount**	**Amount at the best limit just before the aggressive order**
Partial aggressive orders	11 k €	6 k €	49,5 %	38 %	41 k €
Exact aggressive orders	13 k €	8 k €	46,5 %	48 %	13 k €
N-limit aggressive orders	43 k €	22 k €	4 %	14 %	16 k €

Partial and exact aggressive orders constitute the majority of aggressive orders (96%). Individual partial aggressive orders consume a volume (11 k €) almost equal to that consumed by exact aggressive orders (13 k €). N-limit aggressive orders consume an amount clearly more significant than the other aggressive orders (43 k €).

It is important to point out that aggressive orders, and especially exact (resp. n-limit) aggressive orders happen upon particular conditions, when the quantity at the best limit is significantly less than the average quantity at this limit over all the events: 13 k € (resp. 16 k €) just before the exact (resp. n-limit) aggressive orders, versus 57 k € on average. Furthermore, an analysis presented in the appendix shows that market participants submit exact aggressive orders when the LOB is significantly imbalanced[17].

We note that for some members, the majority of their aggressive orders are purposely partial, exact or n-limit. For instance, 96% of the orders of one agency member are partial, 81% of the orders of one liquidity provider member are exact and 75% of the orders of one agency member and one proprietary member are n-limit.

3.3. MEASURE OF POTENTIAM PROFITS ACCORDING TO THE AGGRESSIVE ORDER NATURE

We recall that the potential profit is averaged over all aggressive orders, weighted by the quantity of each traded aggressive orders. In what follows, we define the price impact as the price evolution at a given point in time after the aggressive transaction.

The graph below shows the price evolution around the aggressive order with respect to its nature.

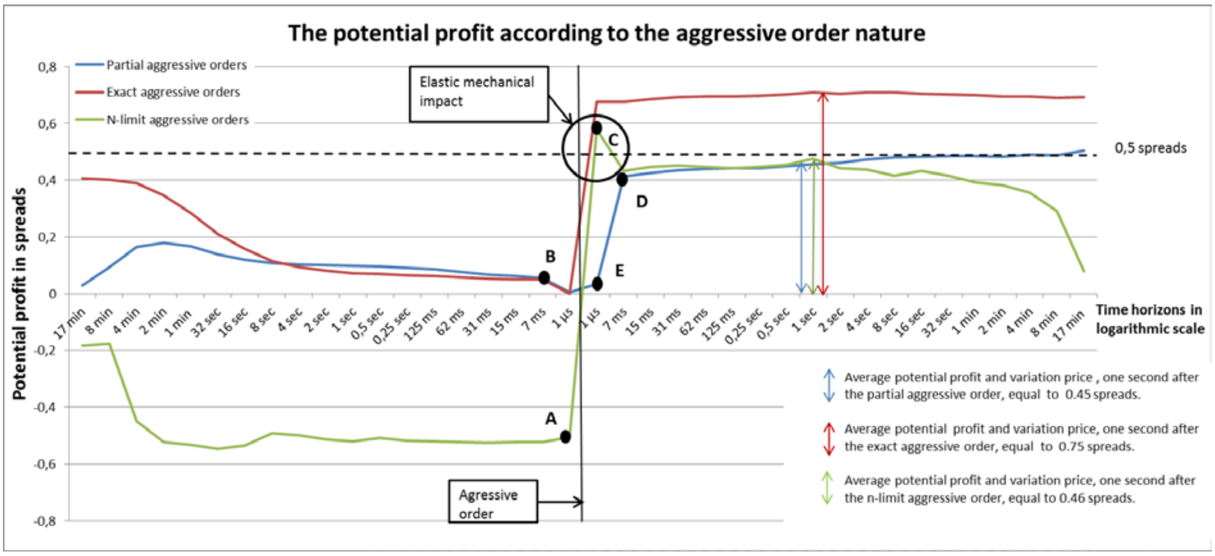

Graph interpretation:

Given the granularity of the available data, no event other than the transaction can take place between one microsecond before and one microsecond after the aggressive order[18].

[17] See in appendix the comparison between the imbalance value reached before a partial aggressive order and the one reached before an exact aggressive order.

[18] There is one exception in the case of a self-trade prevention order that takes place simultaneously with the transaction. Self-trade prevention orders guarantee participants not to consume their own orders: when an aggressive order hits or lifts a passive order issued by the same member, the passive order is automatically cancelled simultaneously with the transaction carried out against other passive orders.

At a given point in time before (resp. after) the aggressive order, **a negative value indicates that the participant could have obtained a better price, at least for one security, by trading earlier (resp. later)** in the case of partial and exact aggressive orders. In the case of n-limit aggressive orders, it is logical to obtain a negative value before the aggressive transaction. For instance, the reached value one microsecond before the n-limit aggressive order is equal to - 0.50 spreads (see point A on the graph above). This is because at least one security was bought at a price higher than the previous best ask. Indeed, the trader could have obtained a better price by splitting his orders and not crossing the limit.

A positive value indicates, on the contrary, that the participant has intervened at an opportune moment. For instance, the potential profit 7 milliseconds before the partial and exact order (see point B) is positive, almost equal to 0.05 spreads. This indicates that if the aggressive order took place 7 milliseconds earlier, the price would have been 0.05 spreads more expensive. Another example is given by the potential profit one microsecond after the partial aggressive order, which is equal to 0.41 spreads (see point D). This means that if the trader has waited 7 more milliseconds to trade, he would have paid at least 0.41 spreads higher[19]. This value also provides information about the potential profit that could be made if the participant succeeds in unwinding his positions passively on the exchange, over a given time horizon.

The difference between the potential profit at a given moment after the aggressive order and the potential profit at a given moment before the aggressive order is equal to the price variation at the best limit between these two moments. For instance, the difference between point C and point A, which is equal to 1.08 spreads, indicates that just after the n-limit aggressive order, the price at the best limit changed by 1.08 spreads.

We point out that in the case of partial and exact aggressive orders, the potential profit is equal to the price variation of the best limit. This is in contrast to n-limit aggressive orders for which the potential profit is significantly lower than the price variation. For instance, one microsecond after the aggressive order, the potential profit of n-limit aggressive orders is equal to 0.58 spreads and the price variation is equal to 1.08 spreads.

As expected, one microsecond after the partial aggressive order, the price is almost the same[20]. Indeed, the price should not change because the next event that can happen in the LOB is almost always higher than one microsecond. For instance, the first percentile of two successive aggressive orders is equal to 13 microseconds. The mechanical impact stands out just after the aggressive order: 1 microsecond after exact and n-limit aggressive orders, the price varies respectively from 0.68 spreads and from 0.58 spreads, due to the mechanical effect.

The potential profit due to exact aggressive orders is higher than that of partial[21] ones, on all time horizons. Their impacts are permanent within 17 minutes. This indicates that when a limit in the LOB is exactly consumed, market participants tend to not submit new orders replacing the executed ones. This is in contrast to n-limit aggressive orders whose impact is elastic (by this we mean that the price tends to go back to its initial value). The

---

[19] This interpretation of the price evolution *after* the aggressive order is valid as long as we consider that the price variation would have taken place in all cases: if the trader behind the aggressive order did not submit his order, another participant would have done it.

[20] One microsecond after the partial aggressive orders (see point E), the price should theoretically not change. On the graph above, the slight price variation is due to a particular order functionality: self-trade prevention orders that guarantee to participants to not consume their own orders: when an aggressive order hits a passive order belonging to the same member, the passive order is automatically cancelled simultaneously with the transaction. If the limit orders of a participant constitute x% of the total volume at the best limit, and this latter submits an aggressive order equal to (1-x%) of the volume, theoretically he consumes less than the quantity present at the best limit, but practically, his limit orders will be cancelled automatically, which brings about the price change.

[21] See in appendix a distinction between partial aggressive orders themselves, according to the consumed share at the best limit.

price impact starts to diminish after the aggressive order. 7 milliseconds after the n-limit aggressive order, the impact decreases by 0.14 spreads. The remaining impact starts to decrease one second after the aggressive order in a way that the potential profit of n-limit aggressive orders becomes lower than that of partial aggressive orders (at a one second time horizon). Following to n-limit aggressive orders, market participants tend to submit new orders replacing the executed ones. Over a 17 minutes time horizon, the remaining impact of n-limit aggressive orders is equal to that of exact aggressive orders, and their potential profit becomes close to 0.

In the following, n-limit aggressive orders are excluded because they do not bring new effective information into the market, knowing that their impact diminishes gradually. Furthermore, they represent only 4% of the total number of aggressive orders.

The potential profit due to exact aggressive orders is higher than that of partial one, over all time horizons.

Following exact aggressive orders, market participants tend not to submit new orders replacing the consumed ones. Consequently, the price impact looks permanent.

After n-limit aggressive orders, market participants tend to submit new orders replacing the consumed ones. Consequently, the price impact attenuates gradually: starting one second after the aggressive order, its potential profit becomes lower than that of partial one, and on a 17 minutes time horizon, the remaining mechanical impact of n-limit aggressive orders is equal to that of exact aggressive orders, but contrary to these latter, the potential profit is almost close to zero on this time horizon.

In the sequel, we compare this measure over the different market participants with same nature of orders.

4. WHO TRADES OPPORTUNELY ?

4.1. WHAT ACTIVITIES ARE THE MOST PROFITABLE?

As already mentioned in paragraph 2.1, members flagged as SLPs (Supplemental Liquidity Providers) are engaged in a market making programme and are considered as HFTs given the programme requirements set by Euronext. Furthermore, all SLPs are identified as HFTs based on the AMF classification method cited previously.

The following table presents the distribution of partial and exact aggressive orders according to the trading capacities:

Trading Capacity	Percentage of partial aggressive orders	Percentage of exact aggressive orders
Agency	27%	16%
SLP	39%	63%
Proprietary	31%	21%
RMO	3%	0%

The two graphs below present respectively the price evolution around partial and exact aggressive orders, according to their respective trading capacity: agency, proprietary, liquidity provider (SLP[22]) and RMO[23]. We

[22] The market making programme of Euronext Paris, named Supplemental Liquidity Provider (SLP) programme, imposes a market making activity on programme members, including order book presence time at competitive prices. In return, they get favourable pricing and rebates in the form of a maker-taker fees model directly comparable to those of the major competing platforms.

recall that the RMO (Retail Member Organization) and the RLP (Retail Liquidity Provider) are members involved in a programme proposed by Euronext, the Retail Matching Facility, specialized in providing liquidity for retails. RLP members are market makers whose orders under this capacity can only trade in front of RMOs, who transmit retail orders. The latter can trade in front of any participant of the market. As RLPs can only trade against RMO orders stemming from retail participants and are therefore protected against adverse selection risk, they can offer more competitive prices. We note that RMOs have access to a particular LOB which is the LOB accessible for all of market participants, plus to the limit orders issued by RLPs. In the following graphs, the RMO potential profit is plotted according to the regular LOB (without the RLP liquidity).

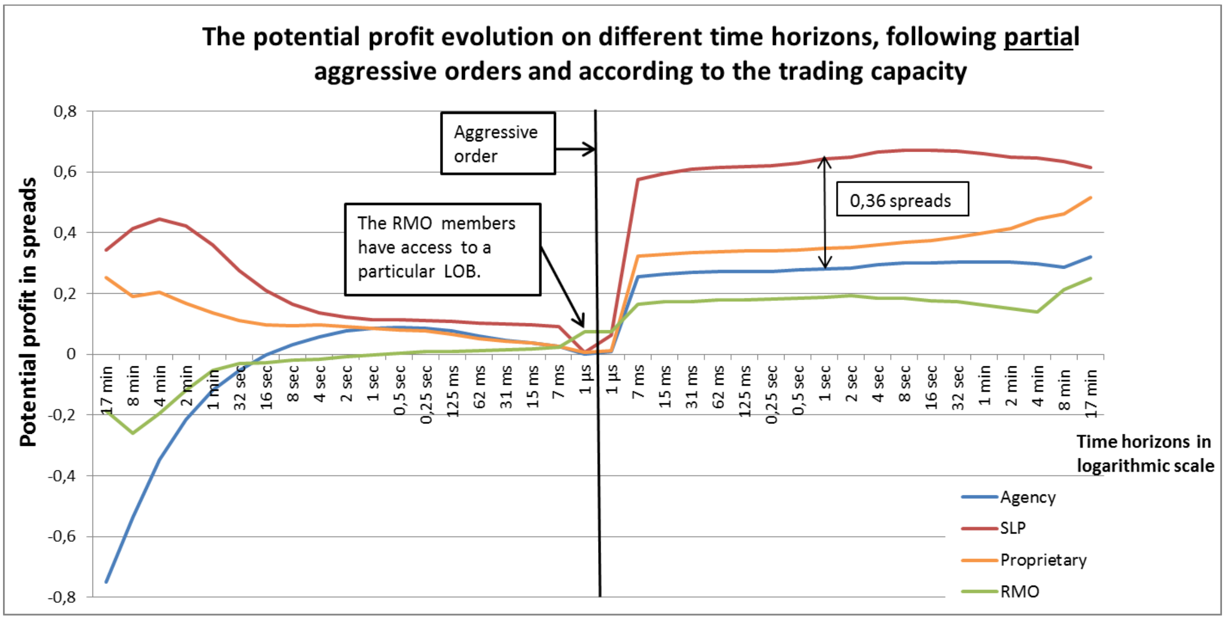

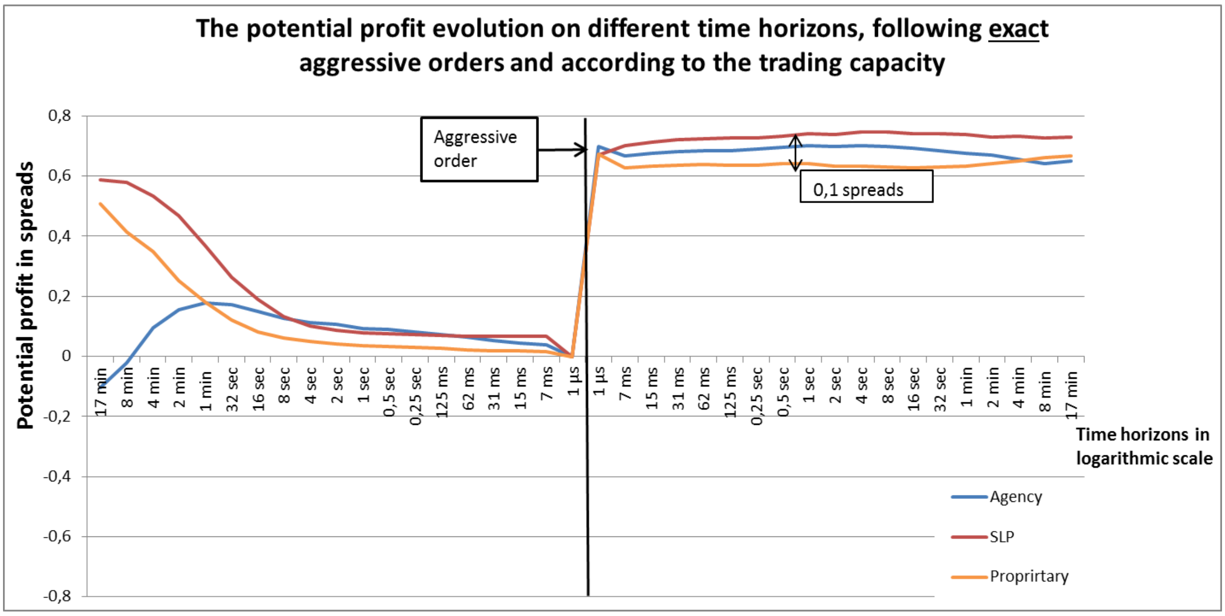

[23] Note that RLP service is not taken into consideration, because the aggressive orders number relative to RLP is little significant. Furthermore, the exact aggressive orders number relative to the RMO service is little significant and therefore RMO is also omitted from the related graphics.

The SLP flow stands out with the (significantly) highest potential profit in the case of partial aggressive orders, on all time horizons[24]. One second after partial aggressive orders, SLPs have a potential profit 0.36 spreads higher than agency participants, and 0.29 spreads higher than proprietary participants.

The results above confirm that HFT market makers (SLPs) make highly profitable use of their aggressive orders: market making is not limited to passive orders only but covers the use of aggressive orders too.

Although the SLPs still obtain a better potential profit than the other trading capacities in the case of exact aggressive orders, the difference in potential profits between the various trading capacities observed is much less significant than in the case of partial aggressive orders: the potential profit of SLPs is only 0.04 spreads higher than agency members and 0.1 spreads higher than proprietary members.

The potential profit 17 minutes before partial aggressive orders is significantly lower than that of exact aggressive orders for each trading capacity, in particular for agency members. This can be explained by the fact that the potential profit 17 minutes before the aggressive orders is inversely proportional to the quantity present at the best limit just before the aggressive order[25]. The following table shows the amount at the best limit before the partial and exact aggressive orders according to each trading capacity.

Trading Capacity	Amount at the best limit just before the partial aggressive order	Amount at the best limit just before the exact aggressive order
Agency	50 k €	19 k €
Proprietary	36 k €	13 k €
SLP	38 k €	12 k €

For all trading capacities, the amount present at the best limit just before exact aggressive orders is significantly higher than that of partial aggressive orders. This explains why the potential profit 17 minutes before partial aggressive orders is significantly lower than that of exact aggressive orders for each trading capacity.

Observation: A few seconds before the aggressive order, RMOs do not seem at an advantage over the best price displayed in the regular LOB; their potential profit is near-zero. In other terms, if RMOs had sent their aggressive orders some seconds before, they would have been able to obtain from the regular LOB the same price obtained a few seconds later due the exclusive liquidity to which they have access. It is only less than 7 milliseconds before the aggressive order that the advantage appears: 1 microsecond before the aggressive order, their potential profit is higher than 0. Despite their access to exclusive liquidity, they are the members with the lowest potential profit over all time horizons after the aggressive order. This is because they are less informed than other participants and because their aggressive orders do not always impact the regular LOB (this is when they consume RLP liquidity). In addition to this, we show in the appendix that RLPs members undergo the lowest adverse selection.

HFTs make a profitable use of their aggressive orders.

The analysis of partial aggressive orders shows significant disparities in potential profits with respect to the various trading capacities. Such aggressive trades are more discriminating than exact aggressive orders. Consequently, only partial aggressive orders will be used in the remainder of this study as a criterion in order

[24] It can be noted that when computing the potential profit, we consider that the position can be unwound at the best price independently of the quantity (acquired by the aggressive order). As a matter of fact, the average amount present at the best price, equal to 57 k €, is significantly higher than the average amount filled by a partial aggressive order equal to 11 k €.
[25] See appendix for further details about the relationship between the potential profit 17 minutes before the aggressive order and the amount at the best limit just before the aggressive order.

to distinguish the participants who trade the most opportunely, by considering the different granularity levels presented in paragraph 2.1.[26]

On average, SLPs are the most profitable market participants. In the sequel, we investigate how this finding varies within each category.

The table below shows the total number of member codes according to each trading capacity and the number of member codes issuing enough/sufficient [27] partial aggressive orders relative to each trading capacity.

Trading capacity	**Number of member codes**	**Number of member codes issuing sufficient partial aggressive orders**
Agency	74	33
Liquidity Provider	17	11
Proprietary	72	24
RMO	23	6
RLP	4	2

The graphs below show respectively, for each trading capacity, the share in number of member codes having a potential profit higher (resp. lower) than the third (resp. first)[28] quartile from 1 microsecond until 17 minutes after the aggressive order. As an illustration, in the first graph below, the value relative to SLP activity 7 milliseconds after the aggressive order is equal to 90%. This means that 90% of the SLPs member codes have a potential profit higher than the third quartile.

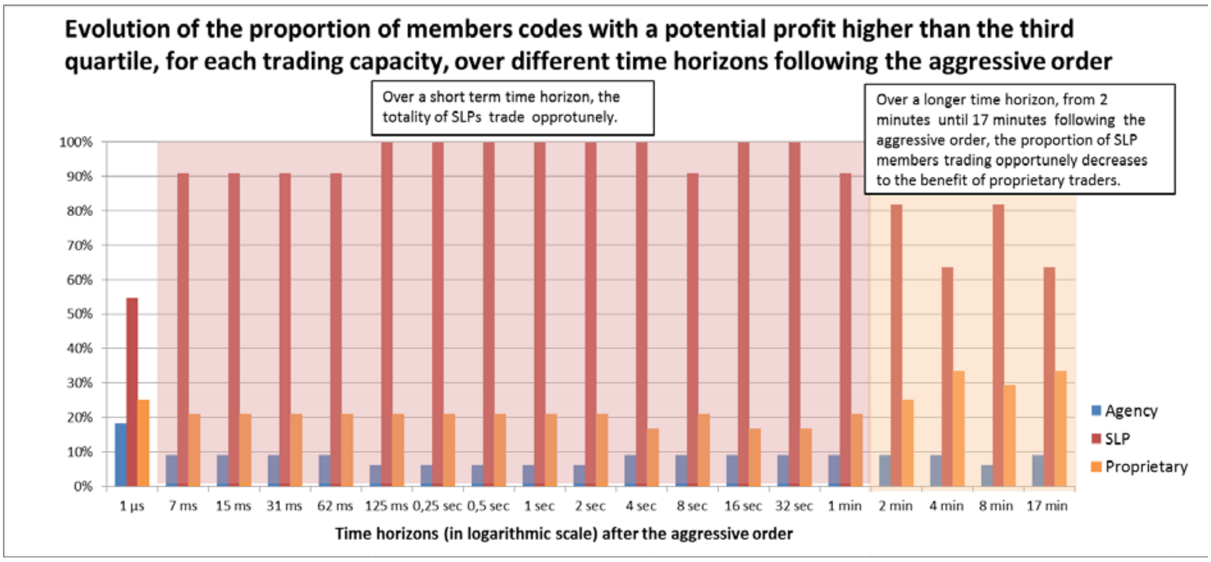

[26] Note that by focusing the analysis on partial aggressive orders exclusively, we were not able to capture members submitting rarely this particular type of aggressive orders.
[27] We consider that the number of aggressive orders is enough/sufficient when there is at least one aggressive order by day and by isin. The study being carried out for/over 66 days, we should have at least 2442 aggressive orders.
[28] The first and the third quartiles are computed based on the potential profits of member codes having enough partial aggressive orders.

Evolution of the proportion of members codes with a potential profit lower than the first quartile, for each trading capacity, over different time horizons following the aggressive order

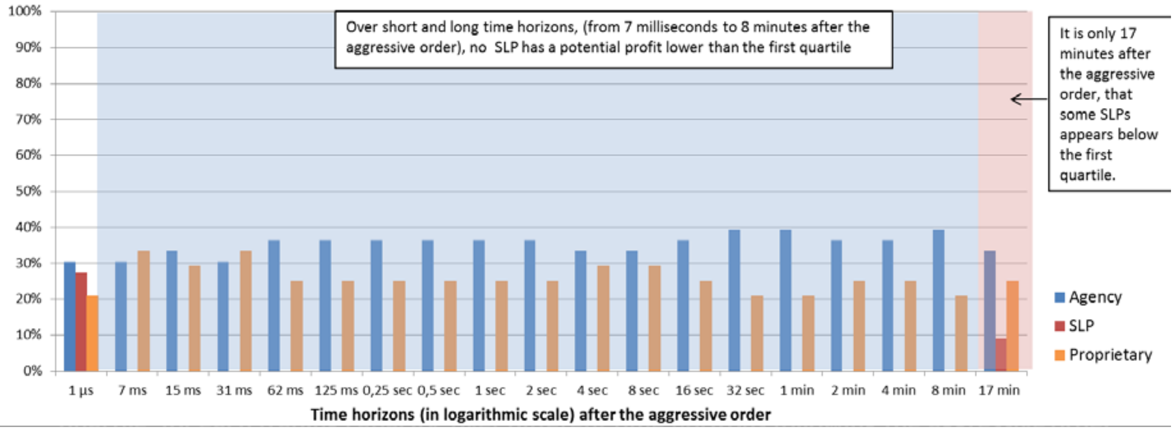

As established in section 3.3, one microsecond after the partial aggressive order, since the price has yet to be impacted, it is not possible to detect the potential profitability. This is the reason for which we do not analyse the potential profit obtained one microsecond after the aggressive order. From seven milliseconds until one minute after the aggressive trade, almost all SLPs have a potential profit higher than the third quartile, and the majority of participants with a potential profit lower than the third quartile are agency traders (35% of them belong to this category).

Over a longer time horizon, from two minutes after the aggressive order, the proportion of SLPs with potential profit higher than the third quartile starts to decrease for the benefit of proprietary traders.

17 minutes after the aggressive trade, some SLPs start appearing in the category whose profit is lower than the first quartile. This is because SLPs do not target long term strategies, high-frequency trading being an activity where participants typically hold positions for very short periods of time/for a very short time spells.

Over a short time horizon, all SLPs are among the members realising the highest potential profits of the market. In the next section, we aim to establish whether HFTs as a whole generate profits as high as SLPs.

4.2. ARE HFTS THE MOST PROFITABLE MEMBERS OF THE MARKET?

We recall that HFTs are identified based on AMF knowledge of market participants and on the estimation of their latency, based on the lifetime of cancelled orders. This means that HFTs studied in this section include SLPs and the rest of HFTs who are not members of the SLP programme. It may be noted that some SLPs member codes are not considered as HFTs but as mixed members.

The table below shows the total number of member codes and the number of member codes issuing enough partial aggressive orders, relative to each category of members.

Category	Number of member codes	Number of member codes issuing sufficient partial aggressive orders
HFT	20	12
MIXED	13	13
nHFT	85	30

Out of the 12 remaining HFTs, 8 are SLPs.

The graphs below show respectively the proportion of member codes (in number) of each class (HFT, non HFT or mixed members), with a potential profit higher (resp. lower) than the third (resp. first) quartile from 1 microsecond until 17 minutes after the aggressive order.

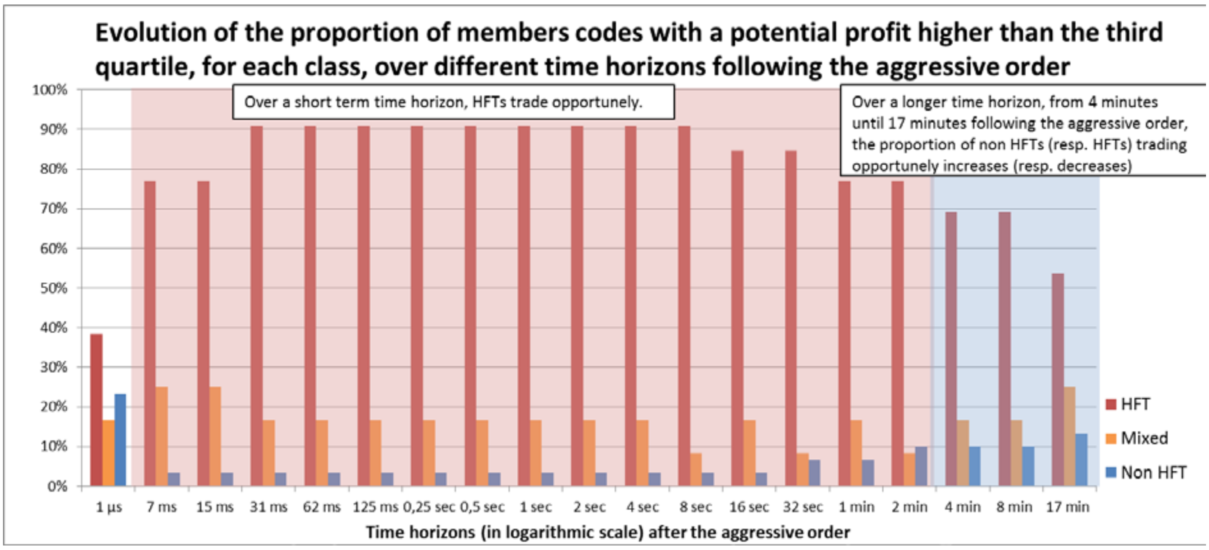

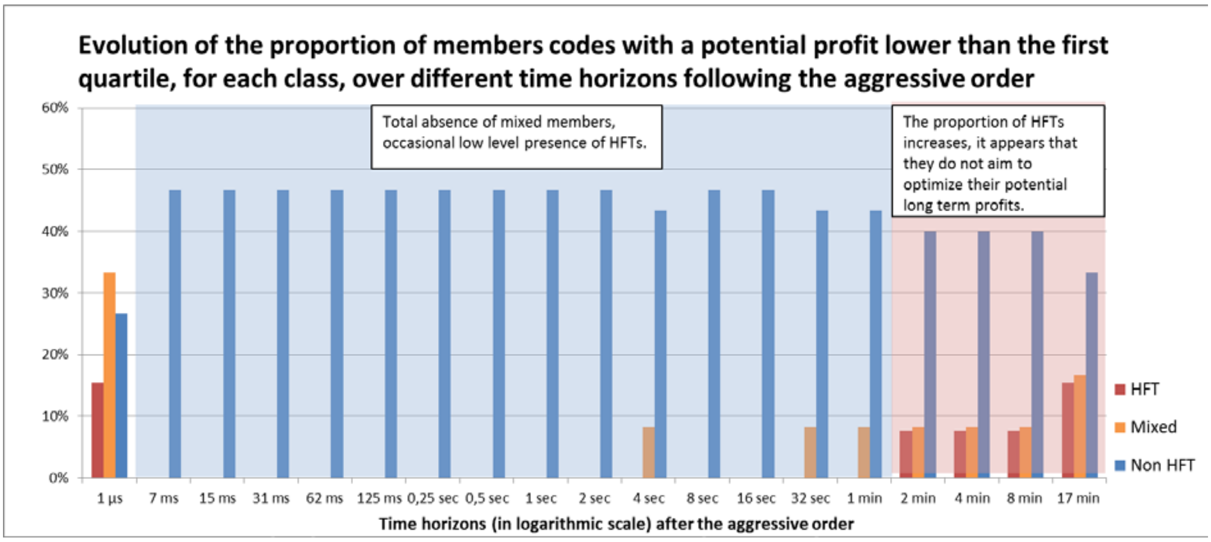

Over a short time horizon (from 7 milliseconds until two minutes approximately after the aggressive trade), the majority of HFTs (between 77% and 92% of them) display a potential profit higher than the third quartile, versus 18% in average[29] for mixed members and only 3.5% in average for non HFTs. Participants with a potential profit lower than the first quartile are almost exclusively non HFTs. Beyond one minute, the presence of HFTs over the third quartile decreases for the benefit of other participants, in particular the mixed members. Simultaneously, the proportion of HFTs, which is lower than the first quartile, starts to increase.

From 31 milliseconds to 4 seconds after the aggressive order, the proportion of HFTs having a potential profit higher than the third quartile is quite constant, equal to 92%. During this time interval, one HFT member code (whose aggressive flows constitute 7% of the total aggressive flows of this member) only does not realise

[29] The average value is computed starting 7 milliseconds until 1 minute after the aggressive order.

potential profits higher than the third quartile. We note that the other member codes of this member are more profitable than 75% of the market participants over the studied time horizon.

On a short term horizon, the potential profit of HFTs is similar to that of SLPs: HFTs almost wholly belong to the members trading the most opportunely.

4.3. DOES THE POTENTIAL PROFIT OF A MEMBER VARY DEPENDING UPON THE TRADING CAPACITY?

In this section, we investigate whether the potential profit of each member code varies depending upon the different trading capacities (agency, proprietary or SLP) it uses.

We found disparities in potential profits in the case of members using different trading capacities at the same time. For instance, member codes belonging to the SLP programme and using simultaneously another trading capacity have different potential profits: the potential profit of the SLP flow comes out always with the highest on a short term horizon. In the majority of cases (except for two member codes), the potential profit of the proprietary flow is higher or equal to that of the agency flow. The example below shows a mixed member code A carrying out simultaneously proprietary and SLP activities.

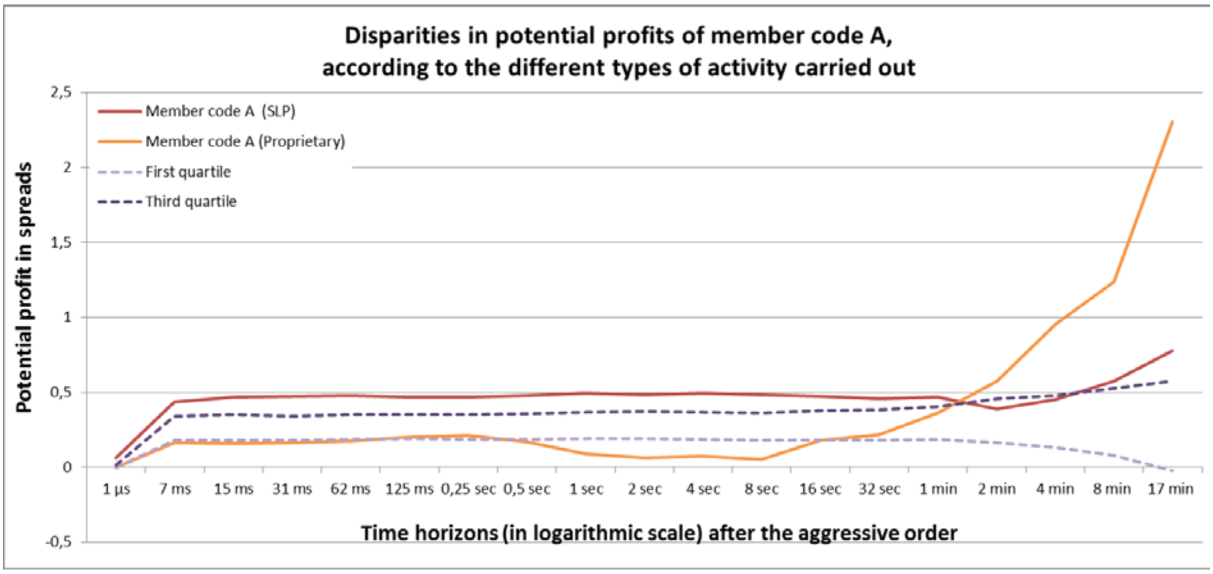

Member code A has different potential profits levels, depending on the trading capacity considered. Until 32 seconds after the aggressive order, the SLP flow has a potential profit higher than the third quartile, while the potential profit of the proprietary flow is equal or lower than the first quartile. The proprietary activity of member code A seems to target a longer term strategy: 2 minutes after the aggressive order, its potential profit becomes higher outperforms the one of SLP. On a 17 minute horizon, it is equal to 2.3 spreads, 3 times higher than the SLP flow potential profit.

Another example is given by dissociating the flows of another mixed member code B carrying out agency, proprietary and SLP activities simultaneously.

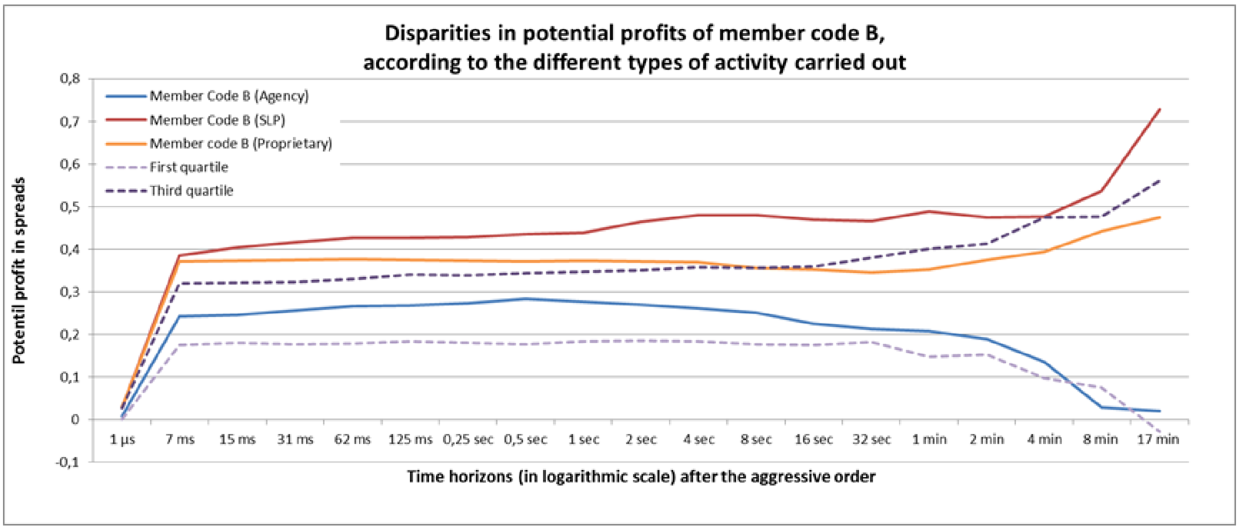

Member code B has different potential profits levels, depending on the trading capacity considered. The potential profit of the SLP flow is the highest: one second after the aggressive trade, the potential profit of the SLP flow (0.43 spreads) is slightly higher than that of the proprietary one (0.37 spreads) and significantly higher than the agency flow (0.27 spreads).

In the light of the fact that SLPs benefit from tariff advantages, it is fair to state that they can take risk more than other members, and in consequence, make less profits than other members. However based on the results obtained, we observe that SLPs are the most profitable members: it seems that their risks are taken wisely.

The dissociation of flows according to the trading capacity allows us to identify the different strategies carried out by a same mixed member: it is possible to distinguish between the high-frequency trading strategies targeting short term potential profits and longer term strategies. For members carrying out SLP and another activity simultaneously, the potential profit of the SLP flow is always higher than that of the other flows. In majority of cases, the potential profit of the proprietary flow is higher or equal to that of the agency flow.

4.4. DOES THE POTENTIAL PROFIT OF A MEMBER VARY ACCORDING TO THE SLE?

In this section, we study the potential profits on a more granular scale using the SLEs of the member, in order to identify disparities in potential profits. An SLE is a Euronext connectivity channel, which members use to convey their orders. Some members may use dedicated SLEs for different strategies.

On the three months studied and among the CAC 40 stocks, there are in total 355 SLEs. 94% of SLEs are used by only one member code, 6% are used by two member codes (belonging to the same member). 73% of SLEs are deployed for only one type of activity (agency, proprietary, SLP, RMO), 26% are deployed for two types of activity and 1% for three types of activities. The number of SLEs belonging to the same member varies between 1 and 33. There are 46 member codes using more than one SLE from a total of 117 member codes.

The analysis based on SLEs does not always provide new information. In some cases, the segmentation of flows belonging to the same trading capacity according to different SLEs does not allow us to identify different trading strategies, and as a consequence nor to spot disparities in potential profits. In fact, potential profits can be homogeneous among the different SLEs. To some extent, this is reassuring, because it proves that the strategies we found in the previous sections are consistent and not random. In some other cases, some disparities in dissociated flows can be observed, showing that some members could use various SLEs to dissociate their strategies[30]. For instance, the graph below shows the potential profit of two flows belonging to the same member code A, belonging to the same trading capacity (proprietary) but passing through two different SLEs.

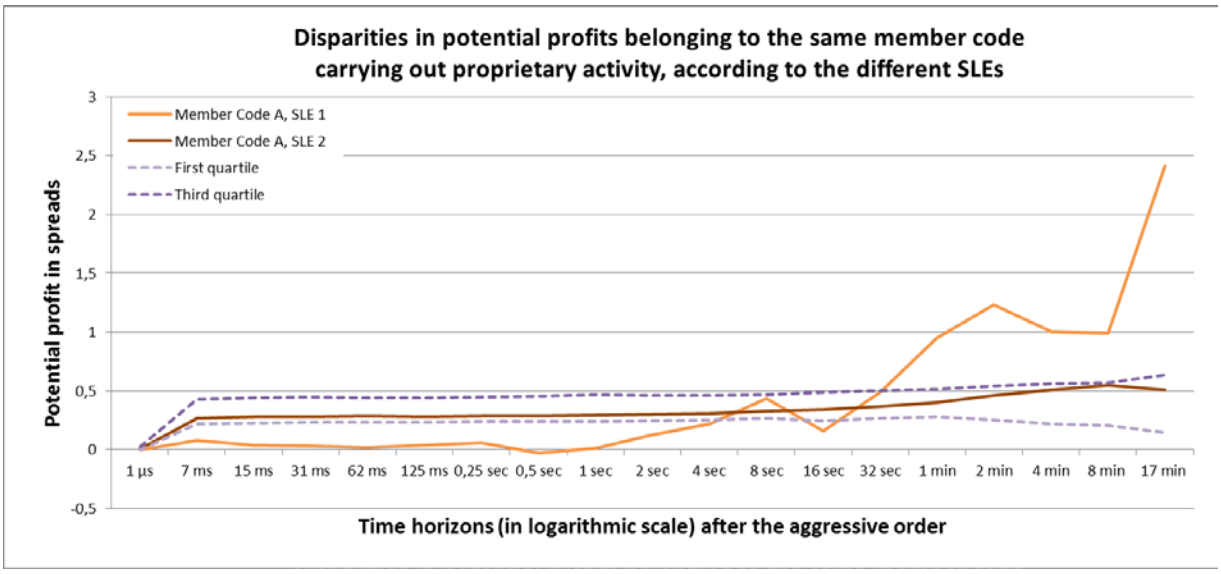

Two distinct proprietary strategies stand out despite being used by the same member code. The first strategy targets a short term potential profit, while the second flow targets a longer term profit: one second after the aggressive order, the potential profit is almost equal to 0 spreads, but around 32 seconds after the aggressive order, the potential profit becomes higher than the third quartile, and reaches, 17 minutes after the aggressive order, a potential profit equal approximately to 2.5 spreads, 5 times higher than the first flow.
The next example shows the potential profits of two flows belonging to the same member code B, under the same trading capacity (agency) but passing through two different SLEs[31].

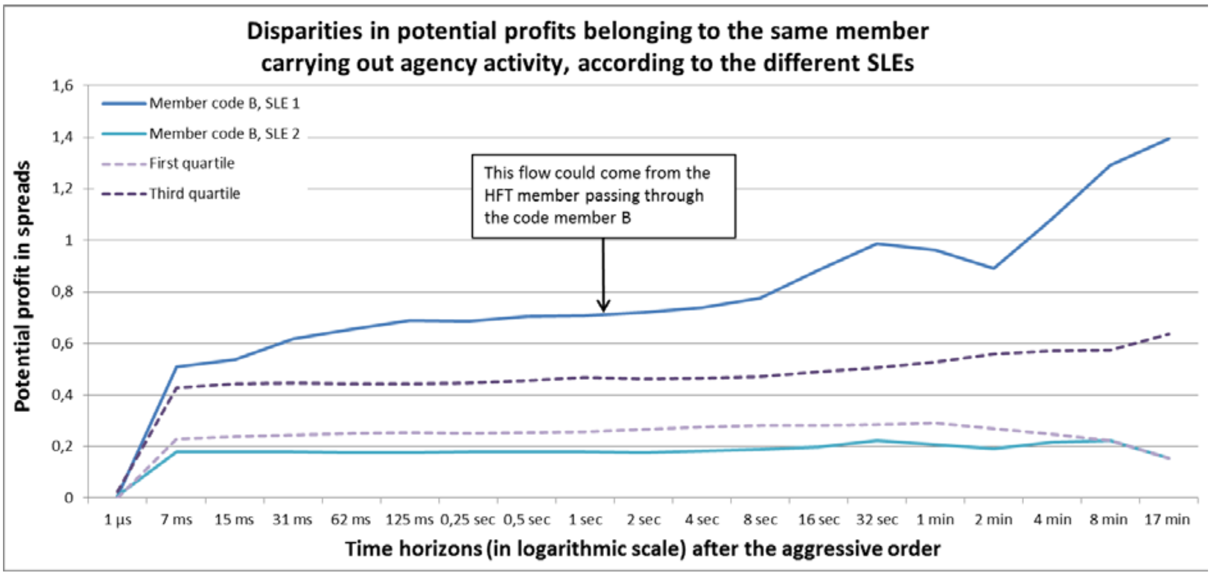

[30] See appendix for further details about separation of aggressive flows and passive flows according to SLEs.
[31] The member code B has more than 2 SLEs, but we choose to show the flows belonging to the 2 SLEs with the highest difference in terms of potential profit.

The member code B is an agency broker, serving as an intermediary for another HFT (among other clients). He seems to have markedly different strategies with respect to the SLEs. This difference in potential profits between the two flows could be interpreted for example by a segmentation by the member code B between the flows belonging to the HFT client and the rest of the flows via the SLEs.

Dissociating the flows issued by the same member code, and belonging to the same trading capacity can in some cases bring up new information concerning the different activities followed by this member code. For instance, by dissociating the flow of a member code carrying out agency activity, it is possible to identify its clients typologies. When dissociating a proprietary flow according to the different SLEs, it is possible to identify different strategies followed by the same member code.

5. THE POTENTIAL PROFIT: MEASURE SERVING DIFFERENT PURPOSES

The measure of the potential profit is not limited to the identification of the participants who trade the most opportunely, but could also serve as a complementary measure to identify HFTs and identify the various market participants' strategies.

5.1. THE POTENTIAL PROFIT: A COMPLEMENTARY MEASURE FOR IDENTIFYING HFTS?

It is frequent to analyse passive orders in order to classify members as HFTs or non-HFTs, by measuring the frequency at which participants update their offers on the market. However, it seems also possible to classify participants by relying on aggressive order potential profit of each member: those realising the higher short term potential profits (one second after the aggressive order) can be considered as HFTs, and those realising the lowest as non-HFTs. The results of this new method are relatively similar to the results of the classification of reference that relies mainly on the analysis of actors' latency. This new method allows us to distinguish more precisely between HFT flows and non-HFT flows belonging to the same member participant if we dissociate flows according to members, trading capacities and SLEs.

The graph below shows the cross of both methods of classification: one relying on latency and one relying on potential profit (the flows are grouped by member code and type of activity).

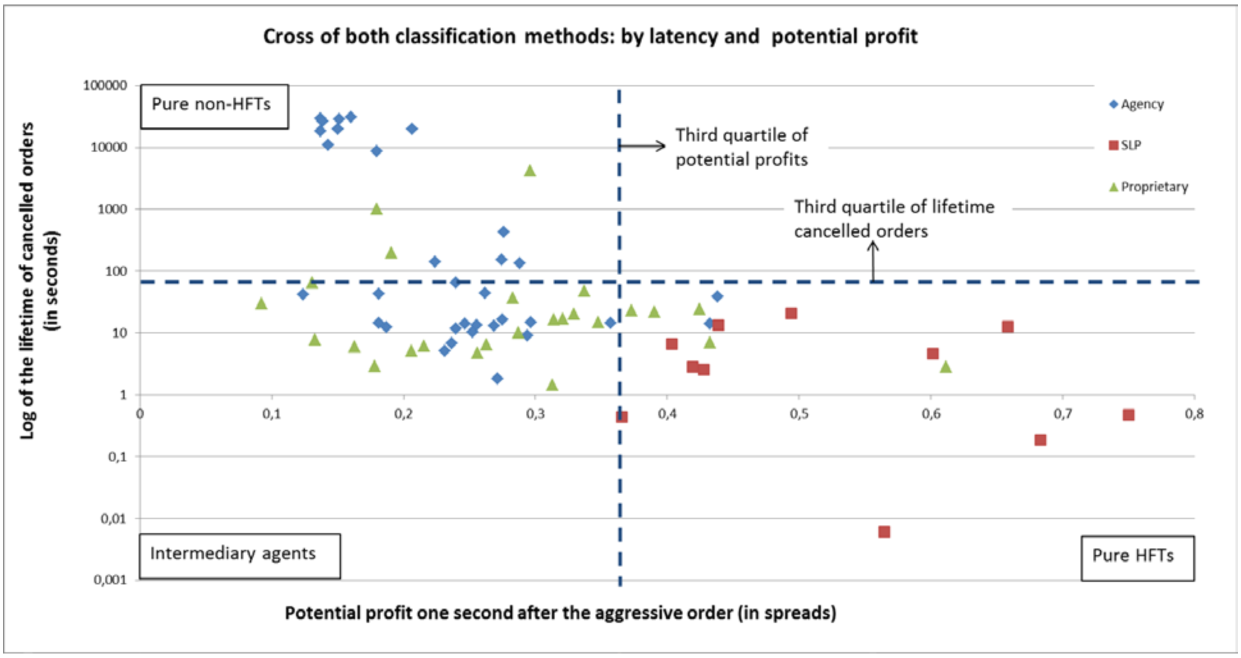

It turns out that crossing both methods allows us to obtain a more complete classification of market participants. Three different categories stand out: pure HFTs characterised by a high short term potential profit and a low lifetime of cancelled orders, pure non-HFTs characterised by a high lifetime of cancelled orders and a low short term potential profit and intermediary agents characterised by a low potential profit and a low lifetime of cancelled orders. We point out that, as expected, no member has high potential profits and high lifetime of cancelled orders at the same time, and that all SLPs figure in the pure HFTs category.

5.2. THE POTENTIAL PROFIT: A TOOL TO IDENTIFY VARIOUS MEMBERS' STRATEGIES?

We draw a distinction between three broad categories of strategies: mean reverting, trend following and another particular strategy described below.
The analysis of the price evolution around the aggressive order can inform us more specifically about the behaviour of each of the participants:

- If the member submits aggressive orders, which direction is opposite to the price evolution (i.e. if he sells while the price is increasing), he likely carries out a **mean reverting strategy**, by going against the price variation.
- If the member submits aggressive orders, which direction is similar to the price evolution (i.e. if he buys while the price is increasing), he likely carries out a **trend following strategy**, by following the price variation.
- If the member submits aggressive orders, just after identifying a new best limit price inserted inside the spread, he likely carries out a **particular strategy** consisting in seizing certain opportunities faster than other market participants.

In the following part of our study, different HFTs strategies are identified and we will observe that some members carry out mean reverting, trend following and another particular strategy (consisting in seizing certain opportunities faster than other market participants) whilst using different member codes.

The mean reverting HFTs

HFTs who follow mean reverting strategies go against the market trend: they buy when the price is decreasing, and sell when the price is increasing.

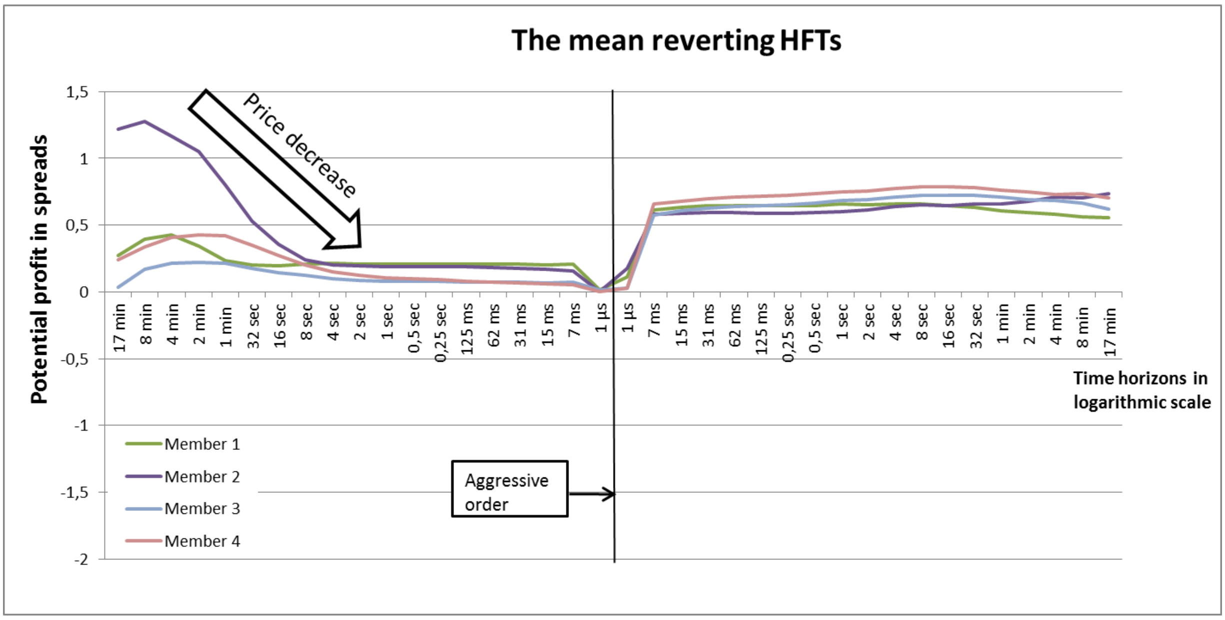

The trend following HFTs

HFTs who follow trend following strategies go with the market trend: they buy when the price is increasing, and sell when the price is decreasing.

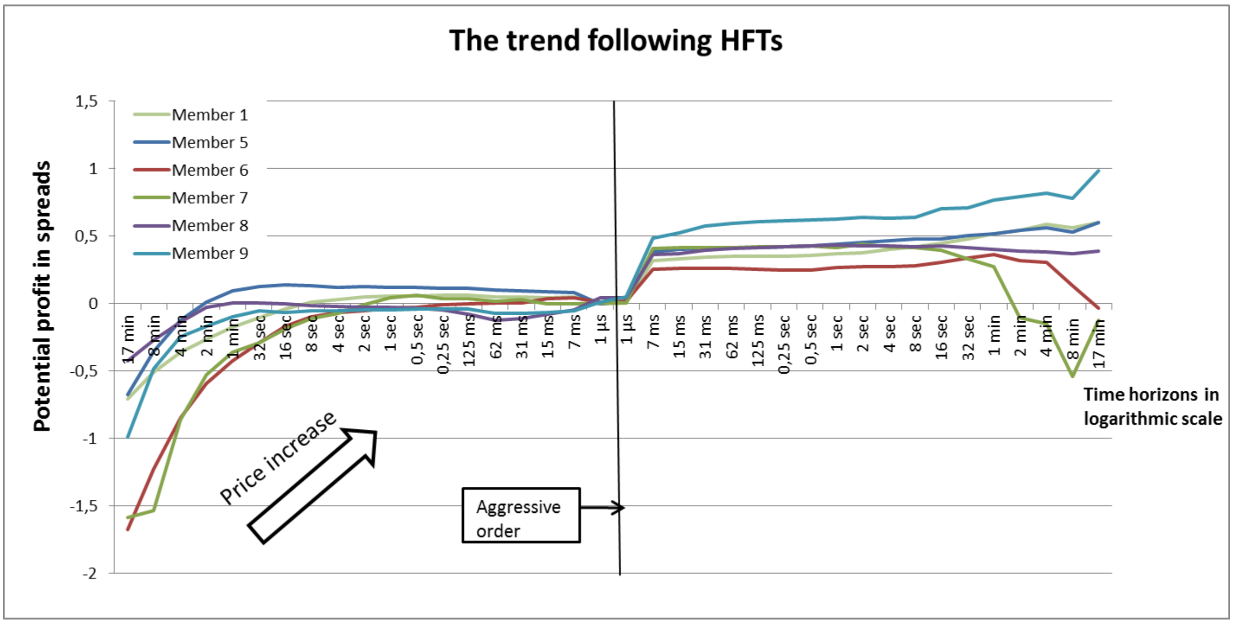

HFTs following a particular strategy (consisting in seizing certain opportunities faster than other market participants)

In the graph below, we focus on participants that intervene when the price is more favourable than what it was a few moments before the aggressive transaction (i.e. when the potential profit is positive a few milliseconds before the aggressive order).

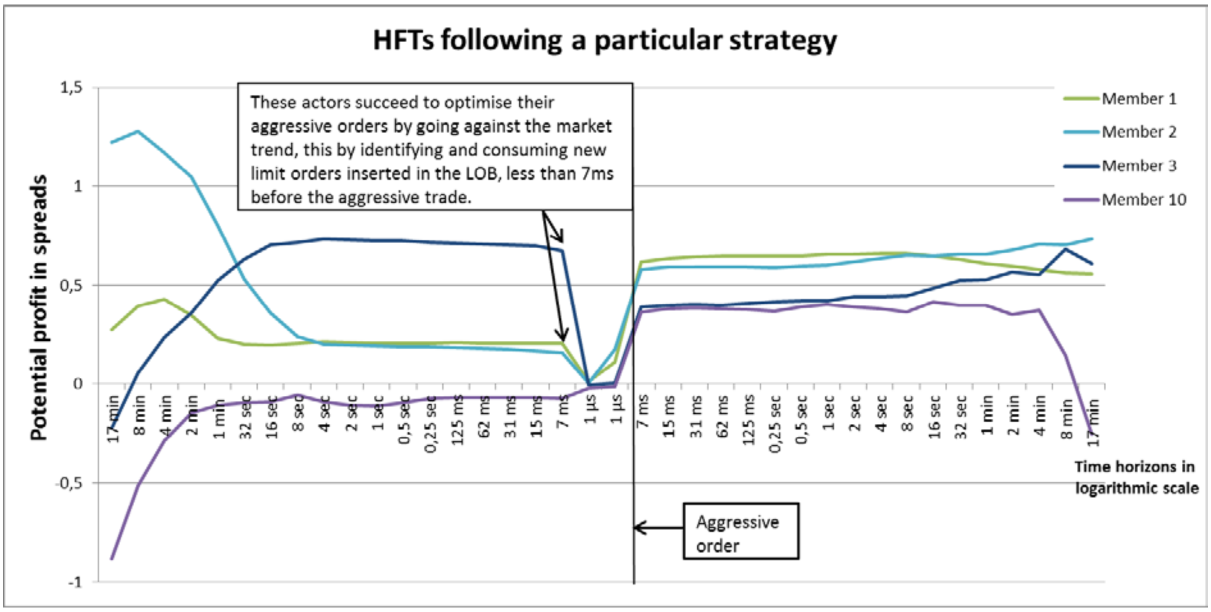

The graph above compares the price evolution around the aggressive trade between three HFTs following a particular strategy consisting in seizing certain opportunities faster than other market participants (member 1, member 2 and member 3) and another HFT (member 10) who does not follow this same particular strategy. The arbitrage configuration clearly stands out: the three members benefit from the insertion of new orders that reduce the spread by 0.6 spreads on average (in the case of member 3) and by 0.2 spreads on average (in the case of members 1 and 2). These situations are likely to occur when the spread is large/wide (equal to several ticks)[32].

Observation: some members carry out distinct strategies simultaneously. For instance, member 1 follows mean reverting, trend following and a particular strategy (consisting in seizing certain opportunities faster than other market participants) at the same time using different member codes. Member 2 and member 3 follow mean reverting and a particular strategy (consisting in seizing certain opportunities faster than other market participants) at the same time, using different member codes. Seen that the analysis is done on a macroscopic scale, we cannot determine if these members use these strategies at the same time, by having for example two different algorithms working at the same time, or if their strategy changes according to the state of the market.

---

[32] The observed results in the graph can be interpreted by the following hypothesis. First, the members could be the fastest to identify and consume the orders recently submitted in the LOB. Second, they could perform instantaneous wash trades that bias the results. Third, they could have a particular behaviour at the beginning and the end of the day that bias their global behaviour among the day. In order to find out which of these hypotheses is the valid one, we performed the same computations, first by excluding all the wash trades of all the participants, and second by excluding all the aggressive trades taking place before 11h and after 16h. The obtained results in both cases are unchanged. We conclude that these three members are the fastest to identify and consume the orders recently submitted in the LOB.

APPENDIX 1: THE IMPLIED VOLATILITY DURING THE STUDIED PERIOD

The implied volatility during the studied period varies little. The three months under study show the lowest implied volatility since 2013.

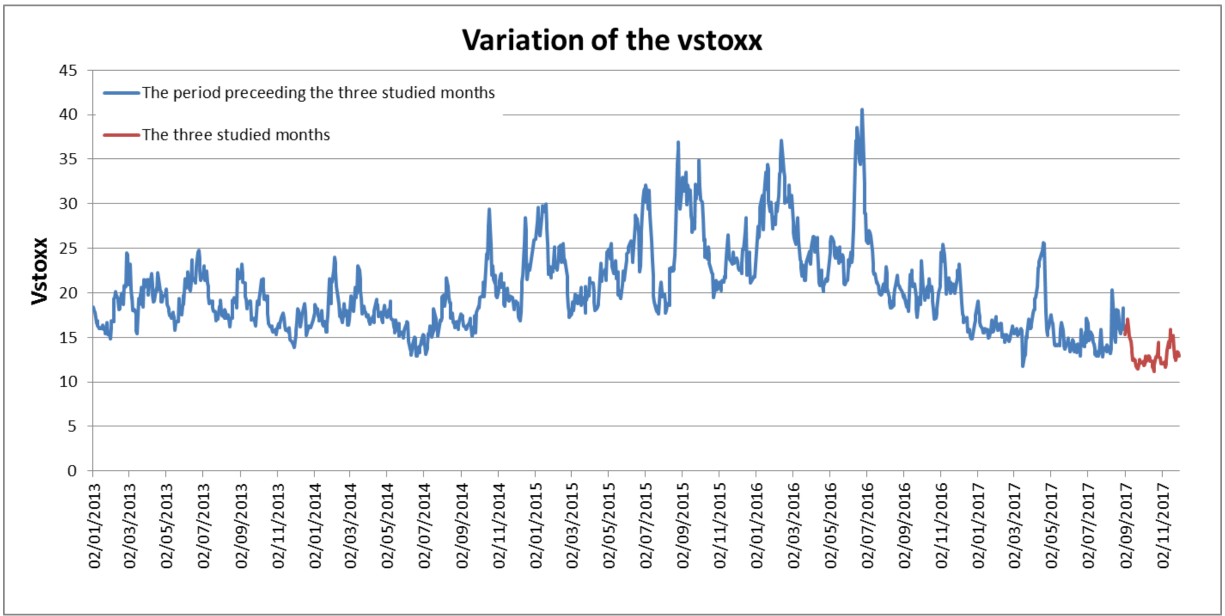

APPENDIX 2: SYMMETRY BETWEEN BUY AND SELL AGGRESSIVE ORDERS

The evolution of the price at the best limit in comparison with the aggressive trade price following a buy aggressive order is quite symmetric to that following a sell aggressive order.

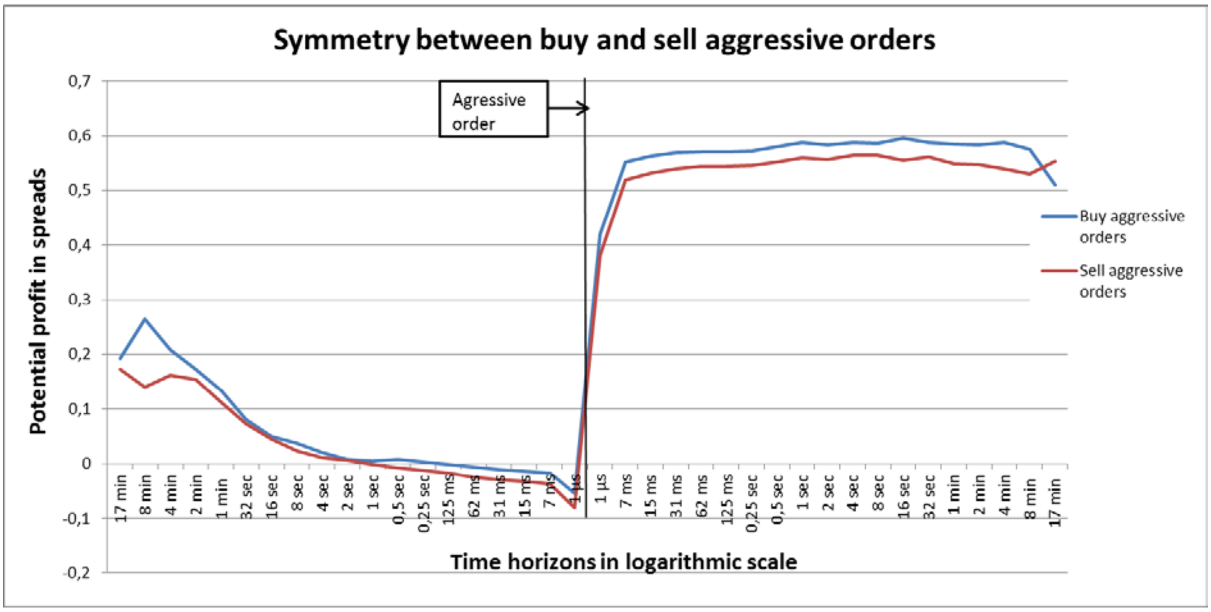

APPENDIX 3: SENSITIVITY WITH RESPECT TO THE UNIT

The measure under examination can be expressed according to different units: ticks, euros, spreads or percentage.

Sensitivity with respect to the unit

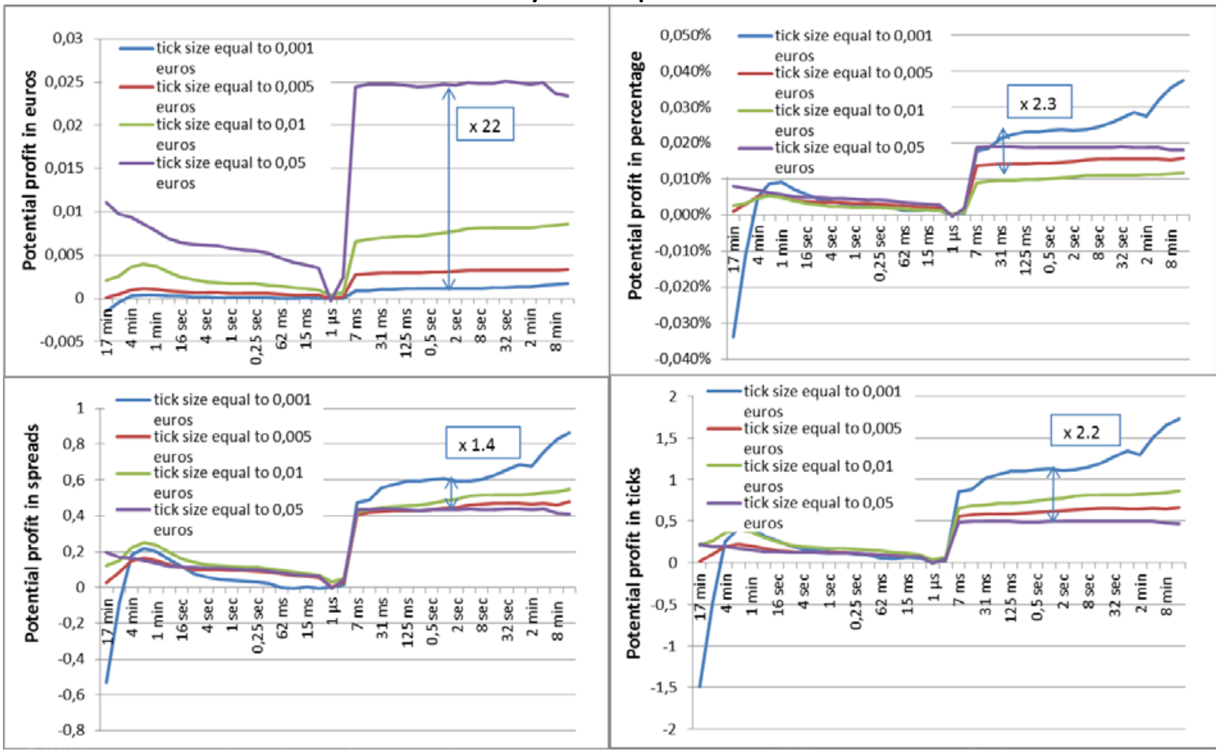

The potential profit expressed in ticks, euros or percentage is sensitive to the tick size: at a one second horizon after the aggressive trade, the maximum ratio between the different tick sizes is equal to 22 when expressed in euros, 2.3 when expressed in percentage and 2.2 when expressed in ticks. We choose the spread as a unit because it provides a basis for comparison among is comparable among all shares (the maximum ratio in this case is only equal to 1.4).

APPENDIX 4: AUTOCORRELATION

The graph below shows the probability that aggressive orders issued using the same trading capacity to succeed on an horizon of 10 aggressive orders.

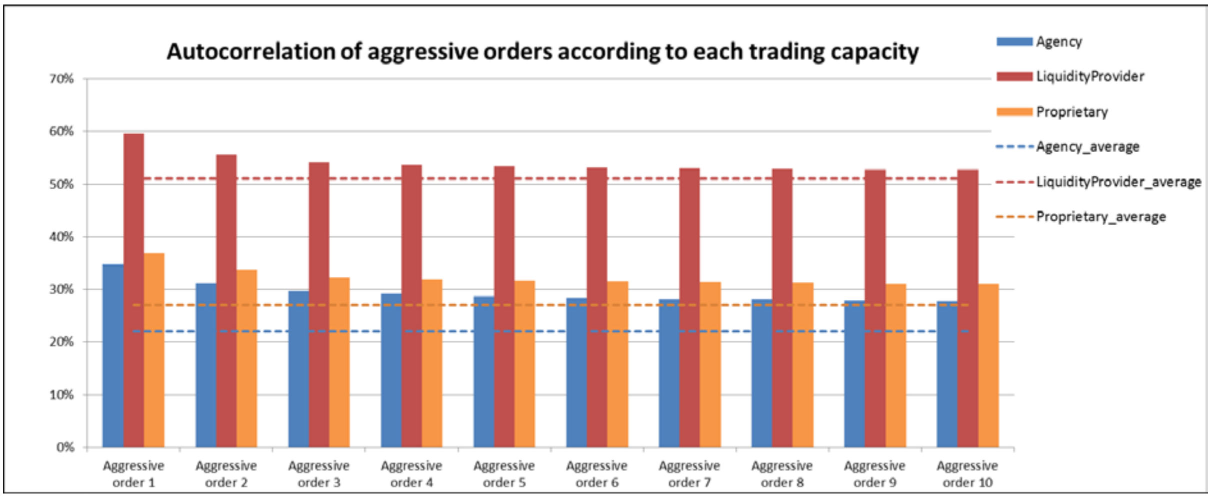

The dark red (resp. dark blue, dark orange) represents the probability that an aggressive order issued by a SLP (resp. agency, proprietary) member follows an initial order issued by a SLP (resp. agency, proprietary) member too. For instance, the probability that the 1st (resp. the 10th) aggressive order, following an aggressive order issued by a SLP, is issued by a SLP member too is equal to 60% (resp. 53%). The red (resp. blue, orange) dotted line represents the probability that a SLP (resp. agency, proprietary) aggressive order takes place in the market at a random time.

In the following, we compute the variation between the probability that an aggressive order issued by a SLP (resp. agency, proprietary) takes place when a previous aggressive order was sent by a SLP (resp. agency, proprietary) compared to the probability that an aggressive order takes place in the market at a random time. The red (resp. blue, orange) plot represents the probability increase that an aggressive order issued by a SLP (resp. proprietary, agency) takes place when the previous (on horizon of 10 successive aggressive orders) aggressive order is also issued by a SLP (resp. proprietary, agency).

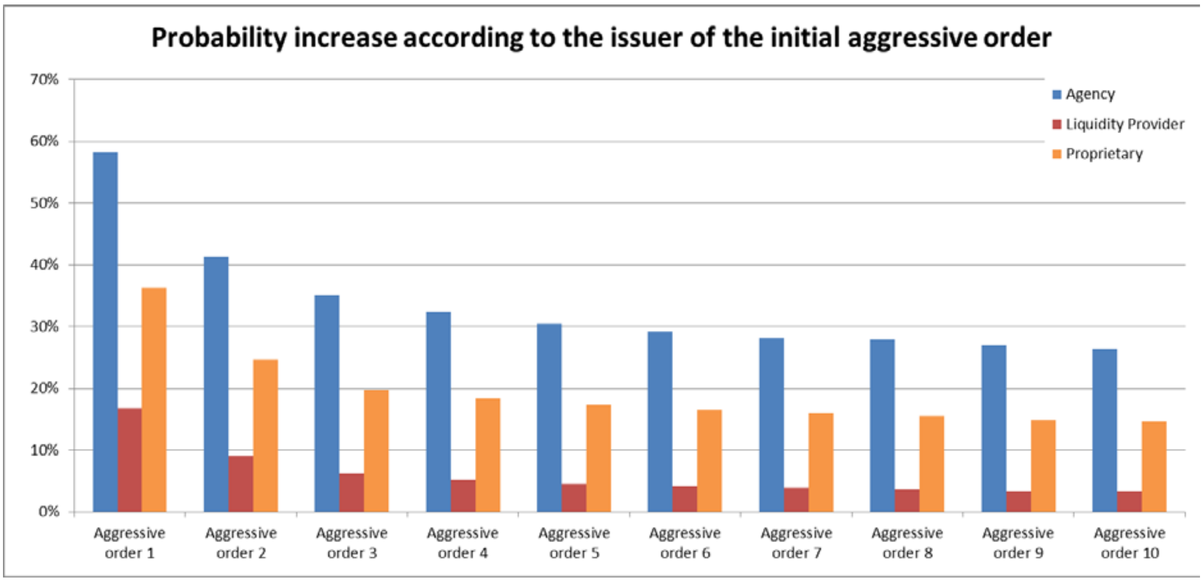

The autocorrelation of SLP aggressive orders is the lowest, compared to the other trading capacities. In fact, the most autocorrelated aggressive orders are those of agency members. When the initial aggressive order is issued by an agency member, the probability that the following aggressive order (aggressive order 1) is issued by an agency member too increases by 58%. This increase is less significant for proprietary members (36%), and significantly lower for SLP members (18%). The autocorrelation of SLP aggressive orders almost disappears for the 10th following aggressive order, while it persists for proprietary (15%) and agency members (27%).

APPENDIX 5: THE IMBALANCE BEFORE A PARTIAL AGGRESSIVE TRADE VERSUS THE IMBALANCE[33] BEFORE AN EXACT AGGRESSIVE TRADE

The value of the imbalance one microsecond before the exact aggressive trade (on average to 27%) is significantly higher than the value reached (on average equal to 3%) before the partial aggressive trades. This result is consistent with the popular rationale saying that the imbalance serves as an indicator for the future evolution of the price. One microsecond after the partial aggressive trade, the imbalance changes mechanically and ephemerally: it increases by 20%. In contrast, after an exact aggressive order, the imbalance reverses its direction ephemerally.

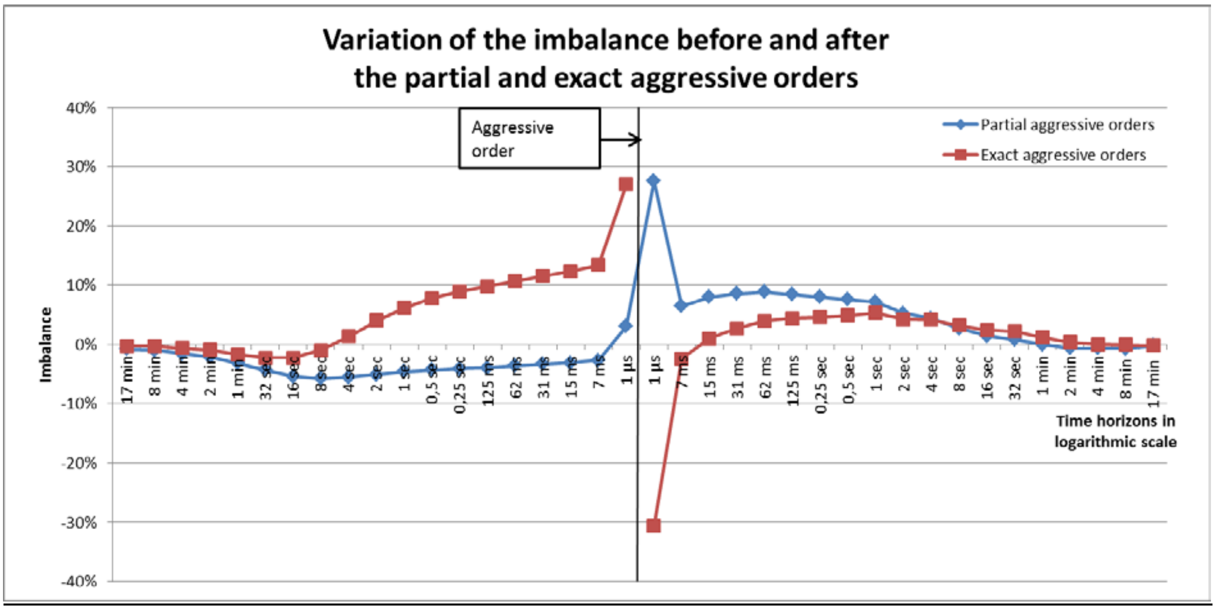

[33] For a buy (resp. sell) aggressive order, the imbalance is computed as the difference between the quantity at the best bid (resp. ask) with respect to the quantity at the best ask (resp. bid), divided by the sum of quantities at the best bid and ask. This ratio is strictly superior to -100% and strictly inferior to 100%. A ratio equal to 0% corresponds to a situation where quantities at the best ask are equal to quantities at the best bid.

AUTORITÉ DES MARCHÉS FINANCIERS
AMF

APPENDIX 6: PRICE EVOLUTION ACCORDING TO THE CONSUMED QUANTITY AT THE BEST LIMIT

The following table shows the proportion of aggressive orders consuming given relative amounts of shares at the best limit (n-limit aggressive orders excluded).

The consumed share with respect to the total quantity at the best limit.	**Proportion**	**Cumulated proportion**
0%-10%	10%	10%
11%-20%	8%	18%
21%-30%	7%	25%
31%-40%	5%	30%
41%-50%	5%	35%
51%-60%	4%	39%
61%-70%	4%	43%
71%-80%	3%	46%
81%-90%	3%	49%
91%-99%	3%	52%
100%	48%	100%

The following graph shows the evolution of the price at the best limit according to the consumed share.

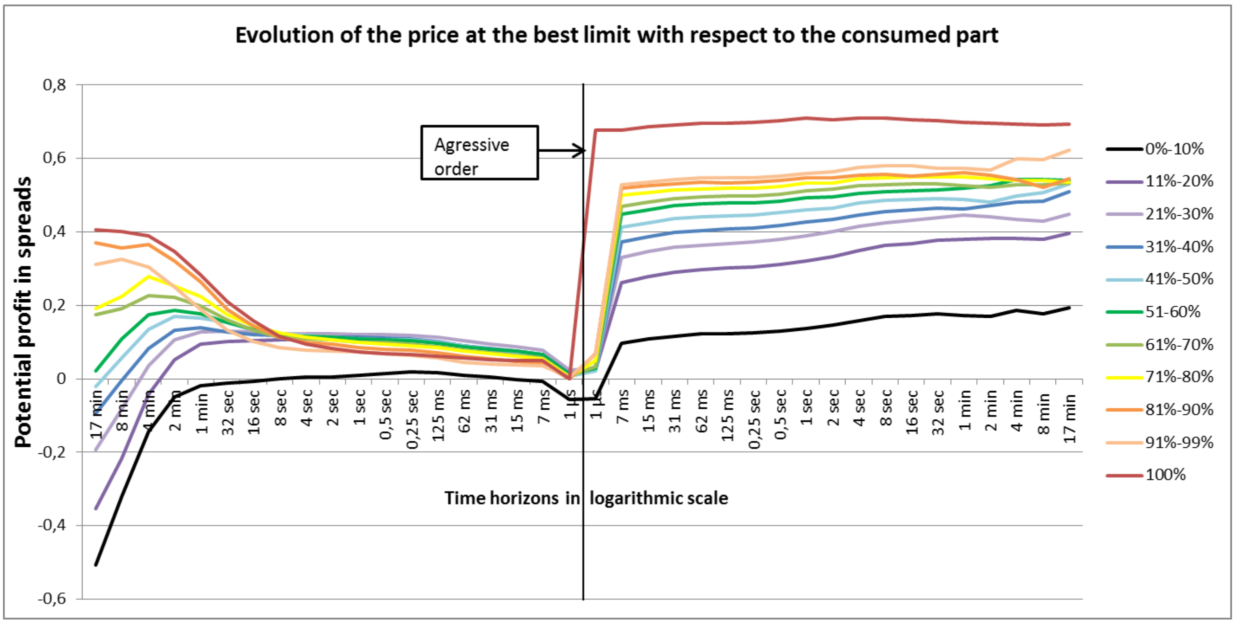

The magnitude of the price evolution is increasing with respect to the consumed share. As an example, the price variation following an aggressive order consuming 10% of the total quantity present at the best limit is significantly lower than the price evolution due to an aggressive order consuming 90% of the quantity present at the best limit. This could be interpreted by the fact that the imbalance created following an aggressive order consuming 90% of the quantity at the best limit is more significant than the one created following an aggressive order consuming only 10%. This likely triggers other aggressive orders or cancellations of limit orders[34].

The following graph shows that the quantity present at the best limit, which is inversely proportional to the consumed part, varies significantly according to this quantitative variable. In contrary, the average traded amount is proportional to the consumed part, and varies less than the quantity at the best limit according to the consumed part.

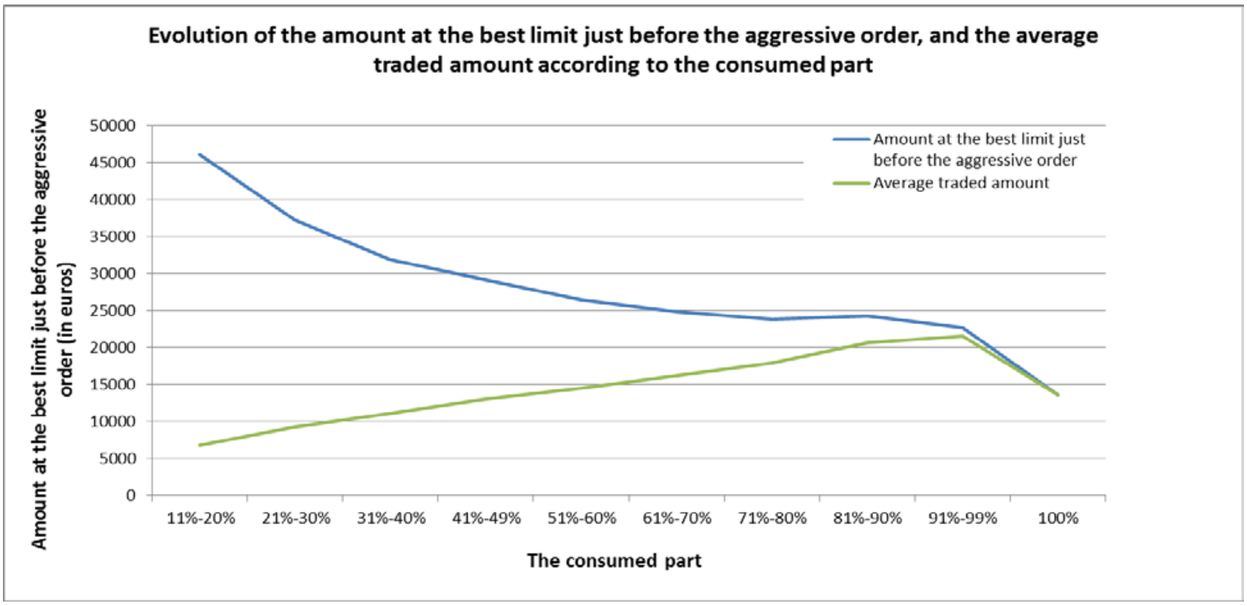

The next graph shows that the potential profit 17 minutes before the aggressive order is an increasing function of the consumed part, which is equivalent to say that it is a decreasing function of the quantity present at the best limit.

[34] A study of the price impact according to the different types of events (aggressive order, limit order insertion and limit order cancellation) can be conducted.

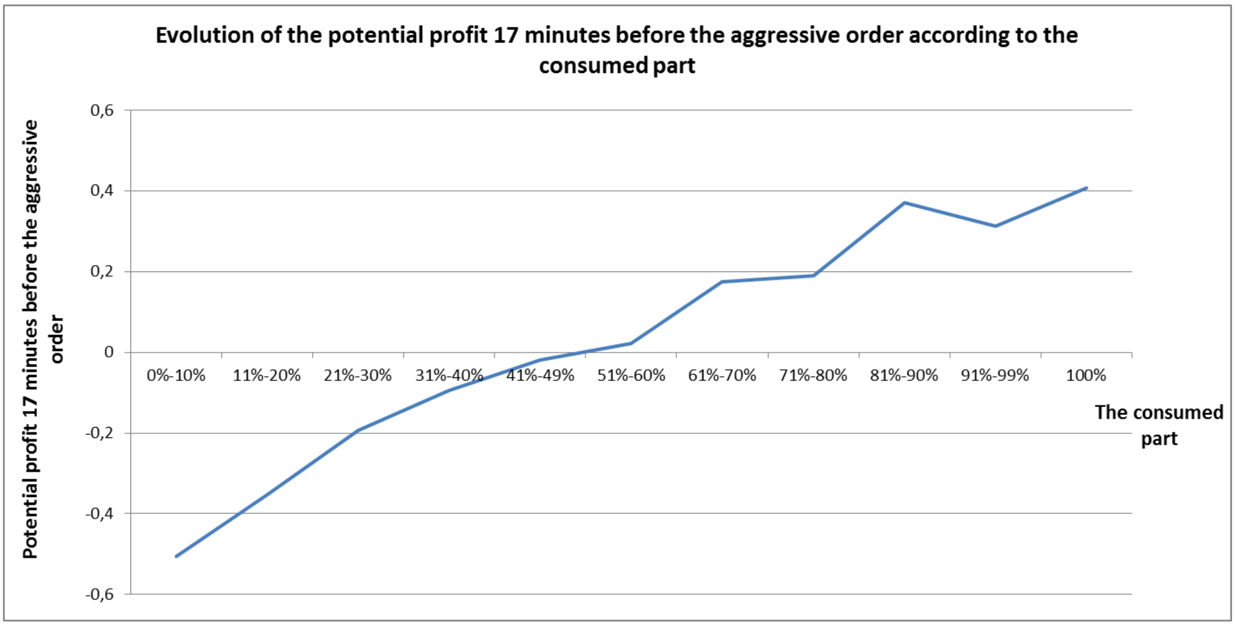

This can be explained by the fact that the quantity present at the best limit at a given point in time, depends on the historical evolution of the price. If the price has been increasing in the last minutes, the quantity present at the best ask tends to increase. In contrary, if the price has been decreasing, the quantity present at the best ask tends do decrease. Following this logic, it is normal to find that the highest the quantity at the best limit (which is equivalent to say: the lowest the consumed part), the lowest the potential profit at 17 minutes before.

APPENDIX 7: ADVERSE SELECTION

We apply the same measure of potential profit on passive orders instead of aggressive orders. The main target is to show that RMOs are the less well informed and that RLPs are the members that undergo the lowest adverse selection in the market.

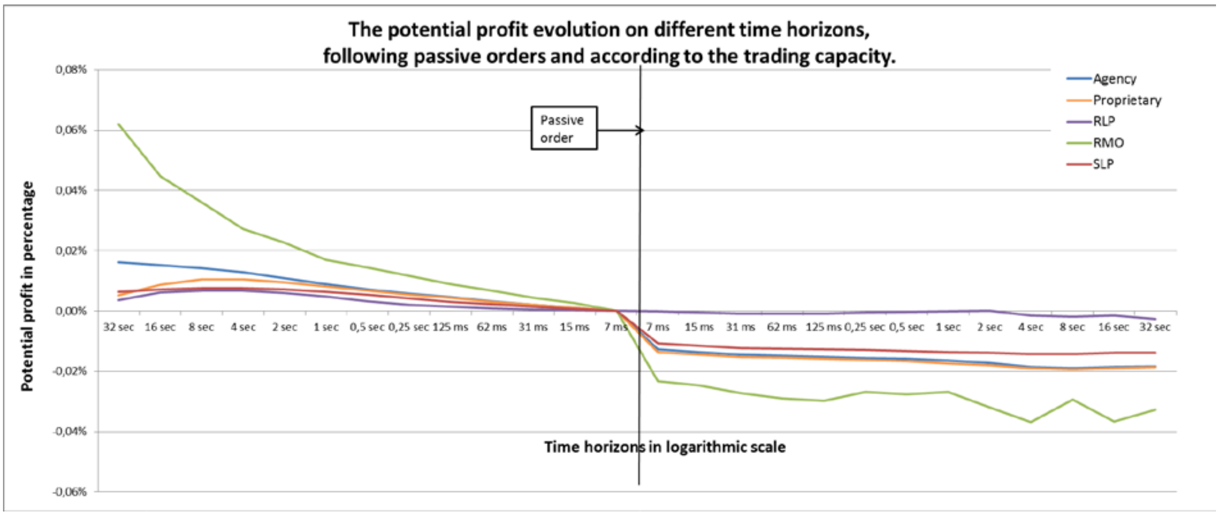

This graph shows that RMOs are the members that undergo the highest adverse selection, while the RLPs undergo the lowest one. These results are consistent with the results found for aggressive orders (in section 4.1). Indeed, RMOs realise the lowest potential profits for aggressive and passive orders.

APPENDIX 8: SEPERATION BETWEEN AGGRESSIVE AND PASSIVE FLOW ACCORDING TO SLES

The analysis of how orders are spread out shows that some member codes separate their aggressive flows from their passive one. For each member code using more than one SLE, we compute the aggressive/passive[35] ratio relative to each of his SLEs. We represent on the graph below the maximum difference between aggressive/passive ratios (each aggressive/passive ratio corresponds to the ratio computed for one SLE) belonging to the same member code. The difference value ranges between 0 and 1. For a given member code, if the maximum difference is equal to 0, this means that all his SLEs are dedicated for the same nature of orders (purely aggressive, purely passive, or mix of both). If the maximum difference is equal to one, this means that the member code uses some SLEs only for aggressive flows only and other SLEs only for passive flows.

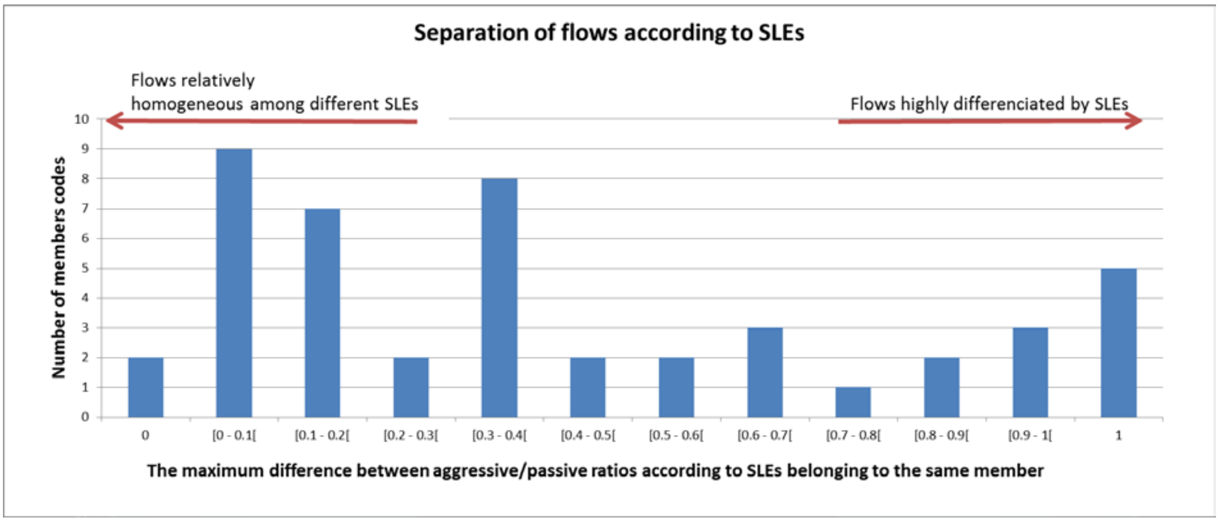

Based on the graph above, we can see that some member codes use their SLEs to separate their aggressive flows from their passive ones

[35] The aggressive/passive ratio is computed as amounts executed aggressively (the order initiating the trade) divided by all amounts traded. This ratio therefore ranges from 0% to 100%; a ratio of 100% corresponds to purely aggressive trades.