Jonathan Tse
+ 44 207 888 2677
Jonathan.tse@credit-suisse.com

Xiang Lin
+ 44 207 888 0974
Xiang.lin@credit-suisse.com

Drew Vincent
+ 44 207 888 4418
Drew.vincent@credit-suisse.com

AES Analysis

High Frequency Trading – Measurement, Detection and Response

Market Commentary **6 December 2012**

Key Points

- While there are a variety of High Frequency Trading strategies – not all of which are bad – the existence of negative HFT strategies has implications for trading and analysis.
- We present a detailed study of a variety of negative HFT strategies – including examples of Quote Stuffing, Layering/Order Book Fade, and Momentum Ignition – to demonstrate what bad HFT "looks like", how often it happens, and how we detect it.
- Among other observations, we find that Quote Stuffing occurs more on MTFs, Order Book fade is more likely on the same venue and less likely cross-venue, and some momentum ignition patterns can cause significant, rapid price moves.
- AES responds to negative HFT with real time detection techniques to target a variety of behaviours. By isolating and identifying different types of HFT, we are able to provide better protections and safeguards for our clients.

What does "bad" HFT look like, how often does it happen, and how do we detect it?

Focussing on the Negative Aspects of HFT

In our previous report High Frequency Trading – The Good, The Bad, and The Regulation, we identified and grouped a variety of High Frequency Trading strategies. We concluded that classifying all HFT as "bad" was too broad a generalisation, as we found evidence of strategies that improved market quality alongside those that did not.

We think it is important to highlight liquidity-enhancing strategies such as market making or statistical arbitrage, which seek to correct short term mispricing. However, this report will focus specifically on strategies which seek to *create* short term mispricing, and how to respond accordingly to this "bad" HFT.

Concrete Examples and Detection Techniques

In this piece we highlight a subset of negative high frequency trading, examining strategies such as: Quote Stuffing, Layering/Order Book Fade and Momentum Ignition. We analyse a number of different aspects of these strategies, providing examples to help demonstrate what they "look" like, as well as broader data statistics on how often they occur and how we detect them.

Exhibits 1 and 2 below provide examples of Quote Stuffing, which is one of the most visually obvious forms of HFT. We will delve into Quote Stuffing in more detail in the next section. We then focus our analysis on Layering/Order Book Fade and Momentum Ignition in subsequent sections. Finally, we highlight the ways in which AES responds to "bad" HFT, protecting our clients and enhancing our strategies.

Exhibit 1: Quote Stuffing: Heineken, 2nd May, 2011

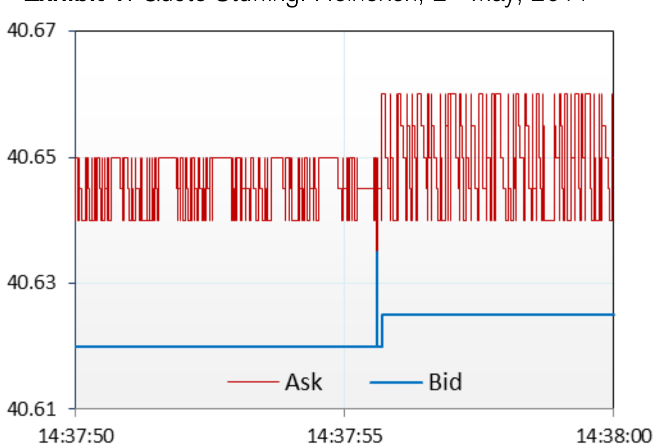

Source: Credit Suisse AES Analysis

Exhibit 2: Quote Stuffing: Telefonica, 10th August, 2012

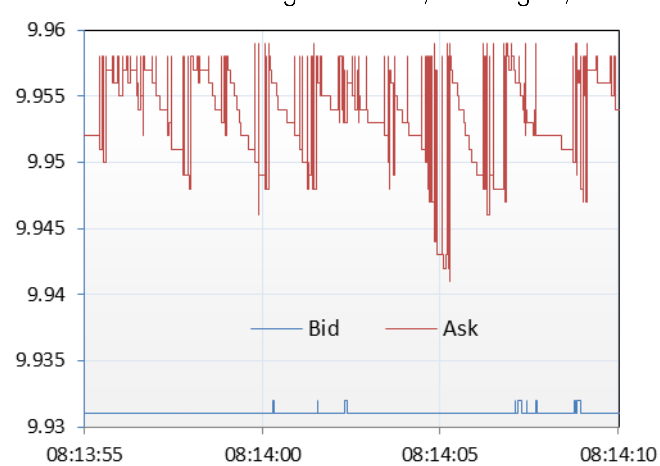

Source: Credit Suisse AES Analysis

What is Quote Stuffing?

Quote stuffing is a strategy that floods the market with huge numbers of orders and cancellations in rapid succession. This creates a large number of new best bids or offers, each potentially lasting mere microseconds.

Why Do It?

This could be used for a number of reasons, including:

- Walking someone into the book: This could game orders that base their pricing entirely on the best bid or best ask.
- Creating false midpoints: One could briefly create a false mid very close to the bid or ask, then trade in the dark (where the mid often serves as a reference) at that price, rather than the "true mid".
- Trying to cause stale pricing, slow market data and suboptimal trading by other market participants: By forcing them to process "false" messages, their trading decisions could be delayed or compromised.

Exhibit 3: Quote Stuffing: Heineken, 2nd May, 2011

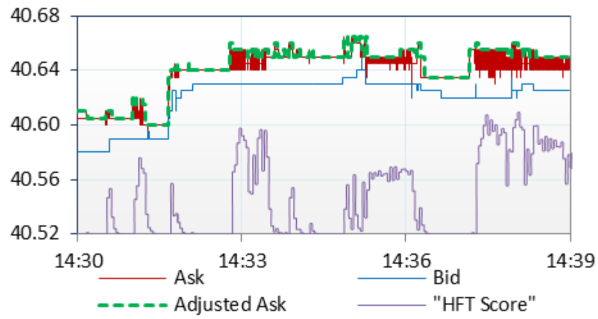

Source: Credit Suisse AES Analysis

Exhibit 4: Duration of Quote Stuffing Events

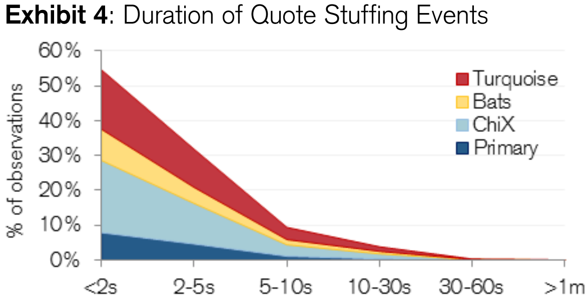

Source: Credit Suisse AES Analysis, STOXX60, Jul – Sep 2012

Exhibit 5: Average Mid-Price move toward quote stuffing (5 seconds post event)

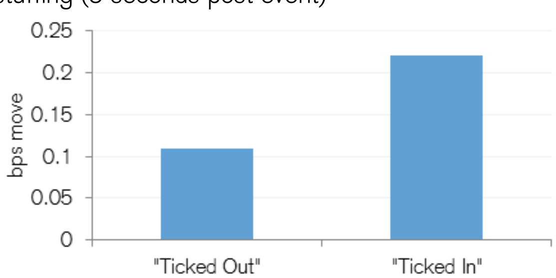

Source: Credit Suisse AES Analysis, STOXX 600 Jul - Sep 2012

High Frequency Quote Stuffing

Catching Quote Stuffing with Burst Detection and Pattern Recognition

The examples on the previous page show quote stuffing with fairly different patterns. The common denominator is the massive number of new orders and cancellations hitting the market in a very short period of time. These bursts are obvious to the human eye, but detecting this across a range of securities in real time – as well as determining the appropriate response – requires some sophistication.

We use techniques adapted from signal processing (including real time burst detection and pattern recognition) to catch quote stuffing and other HFT scenarios. These techniques generate "scores" or "measures" which are updated continually throughout the day on every instrument we trade (more details on these techniques in Appendix 2). We use this information to adapt our trading behaviour accordingly.

Exhibit 3 revisits the Heineken example, now showing a 9 minute window (Exhibit 1 was a 10 second snapshot). Our "HFT score" almost immediately flags this pattern as above the threshold that would trigger behavioural changes in the AES algorithms. Exhibit 3 also presents the adjusted ask, which AES can use to avoid potential downfalls from quote stuffing (more discussion on "Quote Filtering" and other AES protections follow on page 7).

Multiple Times a Day, Across Multiple Stocks

Using the same techniques mentioned above, we analysed the likelihood of quote stuffing across the STOXX600 universe in Q3 2012. We found that the each stock on average experiences high frequency quote stuffing 18.6 times a day, with more than 42% of stocks averaging 10+ events per day.

Mostly Short Lived, But the Long Tail is Important

Unsurprisingly, these events can be quite short lived. In Exhibit 4, we see that the likelihood of events with longer durations is much lower than that of shorter duration events. Indeed, the majority (54.6%) of quote stuffing events (by count) last less than 2 seconds.

However, there is a significant tail of longer-lasting events, which can be several minutes long. While their proportion by count is very low, over 27.9% of the *time* associated with quote stuffing events comes from those lasting 1 minute or more (with over 43.1% due to events lasting 30 seconds or longer). So while most events happen in the blink of an eye, the chance of encountering quote stuffing for over a minute is more than you might expect.

Spreads, Volatility Higher Post Event and Prices Move Too

Although the majority of quote stuffing events only last a short period of time, they can have a significant impact. For instance, we find that average spreads and volatilities are higher in the immediate aftermath of these events. These shifts are over quickly, but they would be taken into account dynamically across all AES strategies to avoid any negative consequences.

On average, the price tends to move *toward* quote stuffing after the event (i.e. the mid-price moves up if quote stuffing occurred on the offer). This holds whether the affected quote finished "ticked in" – narrower than the initial spread – or "ticked out", but is more pronounced when finishing "ticked in" (see Exhibit 5). However, these moves tend to be very small (< 0.23bps).

Exhibit 6: Distribution of Quote Stuffing by Time of Day

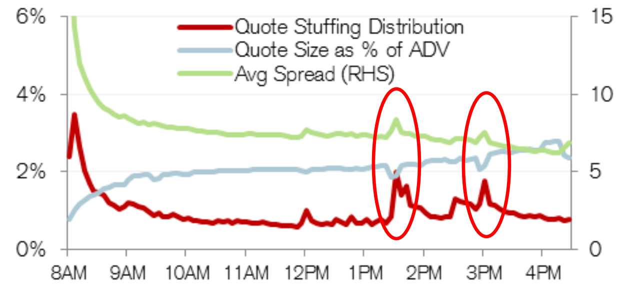

Source: Credit Suisse AES Analysis, STOXX 600, Jul – Sep 2012

Exhibit 7: Time between Quote Stuffing Events

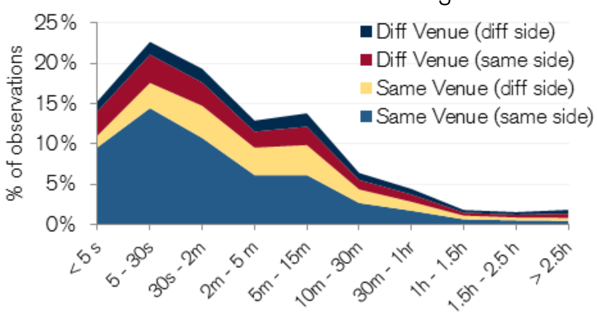

Source: Credit Suisse AES Analysis, STOXX 600 Jul - Sep 2012

Exhibit 8: Kerry Group (Chi-X), 24th October, 2012

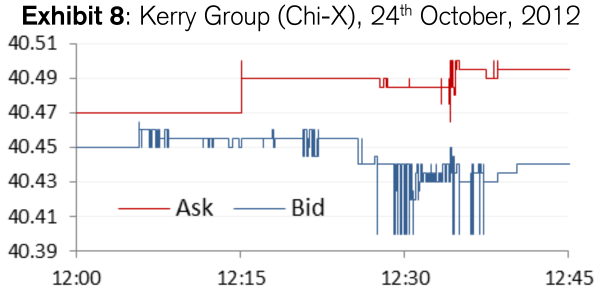

Source: Credit Suisse AES Analysis

Exhibit 9: Distribution of Quote Stuffing Events by Venue

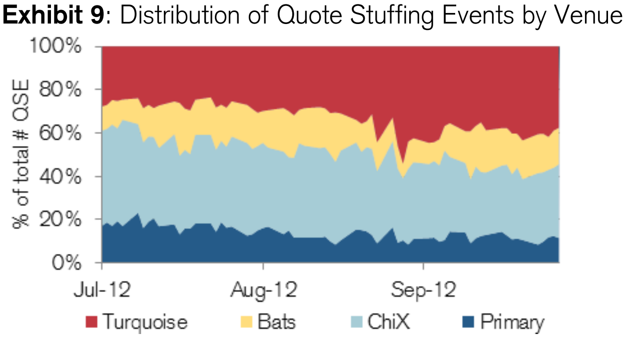

Source: Credit Suisse AES Analysis, STOXX 600 Jul - Sep 2012

Exhibit 10: EBBO Quote Stuffing, Ashmore Group

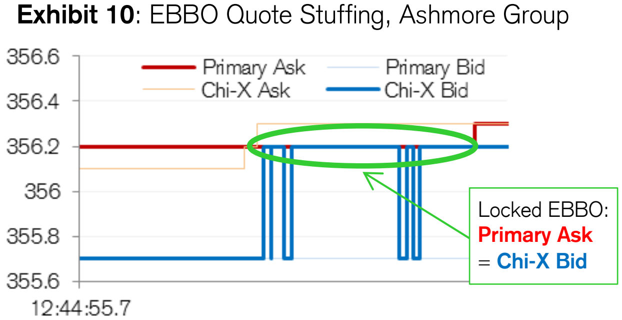

Source: Credit Suisse AES Analysis, 14 Nov 2012

Wider Spreads = More Activity, Repeat Events More Likely

Quote Stuffing is more likely early in the morning, as well as around the time of news announcements (i.e. 1330 and 1500 UK time). Exhibit 6 shows the distribution of quote stuffing (in red) spiking at these times. This coincides with higher average spreads and lower average order book depth, which may provide better opportunities as wider spreads could allow quote stuffing without having to narrow spreads below their "natural" level.

Additionally, a stock which has already experienced quote stuffing has a higher probability of further activity on that same day, with an 82.3% chance of a "repeat" event. The second event occurs on the same venue 73% of the time, and over 70% of "repeats" occur within 5 minutes (see Exhibit 7).

Autos, Banks, Irish stocks see greater quote stuffing

Looking across sectors, we see quote stuffing in the STOXX 600 Autos stocks ~53 times a day (across the 4 venues), and 37 times a day for those in the Banks sector (vs 18.6 times a day overall). We also found Irish stocks more frequently hit, with high frequency quote stuffing events occurring ~95 times a day. This number is dominated by Kerry Group, where we see a huge number of typically short lived events. Exhibit 8 shows a 45 minute snapshot, where multiple distinct clusters of quote stuffing can be seen.

Stuffing Significantly More Prevalent on MTFs

Looking now at which *venue* quote stuffing occurs, we find only 14.1% of events occur on the primary, much lower than Chi-X (37.2%) and Turquoise (32.8%). Exhibit 9 shows this across time, with Turquoise increasing slightly over Q3. It is not entirely clear why MTFs are preferred, but passive rebates – only employed on the MTFs – and newer technology (hence lower latency) could be factors.

Another explanation hinges on a less obvious impact. In addition to slowing down market data feeds, quote stuffing also changes the price that dark pools use as a reference. As there is no trade-through rule in Europe, HFT traders can use multiple order books to create unexpected scenarios.

Implications for Dark Pools, EBBOs and Reference Prices

In particular, dark pools using a synthetic EBBO (consolidated book) for their reference price are at higher risk of being gamed by quote stuffing. Exhibit 10 shows an example in Ashmore Group, where the Primary Bid and Ask (represented by the outer dark red and light blue lines at 356.2 and 355.7) are static, but the Chi-X bid moves (dark blue line). The consolidated EBBO shows a locked book, with the bid equal to the ask at 356.2.

EBBO Pools May Cross Peg-to-Mid Orders at the Touch

This scenario could be exploited in EBBO-referenced dark pools. A gamer could place a sell order in the pool with a 356.2 limit, then place (and rapidly cancel) a Chi-X bid, also at 356.2. Any buy order pegged to mid would trade at the temporary gamed "mid" of 356.2 (as the EBBO bid and offer are both temporarily 356.2), paying the whole spread rather than half.

Crossfinder (Credit Suisse's dark pool) does *not* use the EBBO, preferring to use primary-only data to help minimise the chance of midpoint gaming. Furthermore, when AES detects any quote stuffing, it may add extra protections across its orders (both lit and dark) to further reduce the risk of being gamed, more details of which are discussed later from page 7.

Layering and Order Book Fade

Transient Volume and Unwanted Cancellations

Layering is another frequently cited form of negative HFT. This may take the form of a trader placing a number of sell orders – often at several price points – to give the false impression of strong selling pressure and drive the price down. Then, the trader buys at the cheaper price and cancels the sell orders.

Layering is more viable for high frequency traders. Their speed allows them to mitigate the risk of someone trading against those "false" orders by cancelling immediately in response to any upward moves. This means the buyer gets less than what was displayed on the screen – a common complaint of clients. This can show up in two particular scenarios, discussed next.

What is Layering?

Layering takes the form of a trader placing a number of sell orders – often at several price points – to give the false impression of strong selling pressure and drive the price down. The same holds for a buy.

Why Layer the Book?

By driving the price down, the trader can then buy the stock at an artificially cheap price and trade out when the book reverts.

Price Fade: the Elusive Bid Behind

"Price fade" refers to volume disappearing on a venue as soon as you trade there – e.g. after you buy the 100s, the 101s cancel immediately. While layering is not always the culprit, it undoubtedly adds to the frequency of price fade - an HFT trader at 101 could be cancelling to avoid adverse selection.

Exhibit 11 shows a real example where a trade of 100@29.13 in Legrand SA on Euronext Paris lead to the 1200 shares at 29.135 being cancelled within milliseconds. This behaviour can impact performance and fill rates, particularly for aggressive trading that targets multiple levels of displayed liquidity.

Using tick data, we analysed several markets – again Q3 2012 – to examine how often price fade occurs. We split our analysis into two groups: "full take" - trades that took out the entire price point - and "partial take" - where some volume is left behind. In our definition, "fade" occurs when volume is cancelled after a trade, within one second *and* prior to the next trade.

What is Price Fade?

"Price fade" refers to volume disappearing immediately after a trade, on the same venue.

Why Might it Occur?

One of the reasons why this occurs is that traders cancel orders in response to trades to avoid adverse selection. This is more likely when that trader may not actually intend or need to trade – e.g. in a layering scenario.

Exhibit 11: Price Fade Example, Legrand SA (Paris) September 21st, 2012

Source: Credit Suisse AES Analysis

Price Fade more likely when taking the entire touch

Exhibit 12 shows the likelihood of price fade aggregated across a number of markets[1]. On a "full take" (in grey), the likelihood of price fade is higher compared to a "partial take" (blue), regardless of venue. A full take on the primary results in "fade" 43% of the time, but a partial take leads to cancelled volume only 21% of the time. The difference is smaller on MTFs (38%vs 30%), but the increased likelihood of price fade after full takes still exists. One explanation could be that participants may react more actively to an update in the quote price (vs a *size* update only on a partial take).

Exhibit 12: Likelihood of Price Fade, Across Markets[1]

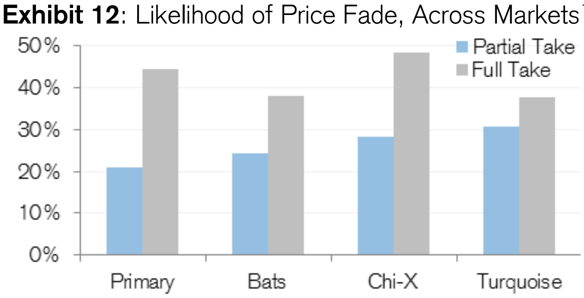

Source: Credit Suisse AES Analysis, Jul - Sep 2012

But Less Frequent When Spreads Are Wider

Exhibit 13 shows the intraday likelihood of price fade across markets, aggregated across venues. In the morning and around economic news releases (1330 and 1500 UK time), the chance of price fade decreases for full takes. This may reflect a view that wider spreads are less susceptible to adverse selection, and more likely to revert. The likelihood of price fade also increases slightly after the US open, suggesting that liquidity from traders who also trade the US – or ramp up when the US opens – may be more transient.

Exhibit 13: Likelihood of Price Fade by Time of Day, Across Markets[1]

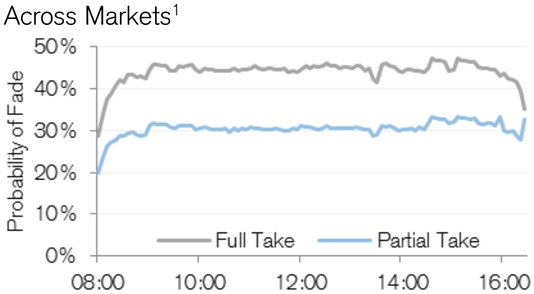

Source: Credit Suisse AES Analysis, Jul - Sep 2012

And more likely now than in previous years

The full/partial take differential holds when looking back in time – Q3 2011 and Q3 2010 show a similar relationship (see Exhibit 14). Interestingly, the likelihood of seeing price fade has significantly increased, especially for "full takes", which could potentially be due to both improved infrastructure across the board, lower latencies, and an increase in colocation.

[1] Copenhagen, Helsinki, London, Madrid, Milan, Oslo and Stockholm. Additional breakdowns are shown in appendix 3 as well as a discussion regarding markets not included, for both Price and Venue fade.

Exhibit 14: Likelihood of Price Fade, Across Markets[1]

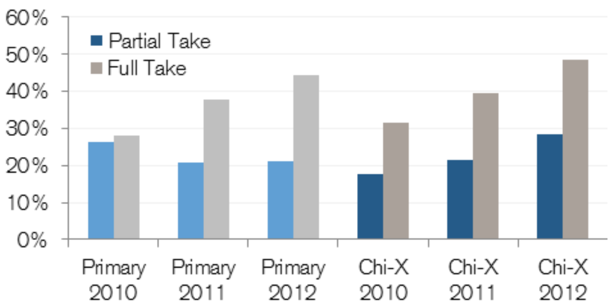

Source: Credit Suisse AES Analysis, Q3 2010, Q3 2011 and Q3 2012

What is Venue Fade?

"Venue fade" refers to volume disappearing immediately after a trade, but on a different venue from the executing venue.

Why Might it Occur?

Similar to Price Fade, this may occur as traders cancel orders in response to trades across venues in order to avoid adverse selection. However, orders posted through SORs may also redistribute volume as executions occur – cancelling volume from other venues in order to repost on the "active" venue.

Exhibit 15: Venue Fade, Unilever Sep 28th, 2012

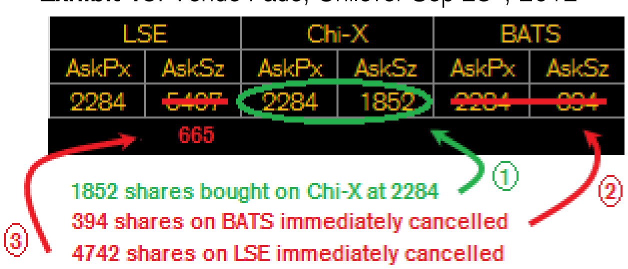

Source: Credit Suisse AES Analysis

Exhibit 16: Likelihood of Venue Fade Following a Trade on the Primary (across markets[1])

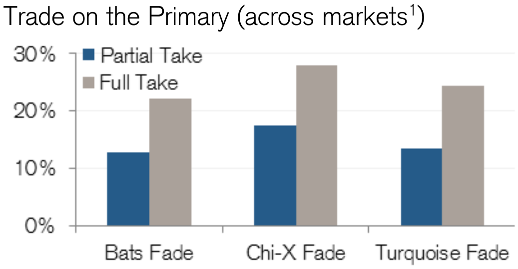

Source: Credit Suisse AES Analysis, Jul - Sep 2012

Exhibit 17: Likelihood of Venue Fade, following a trade on the Primary (across markets[1] – Q3)

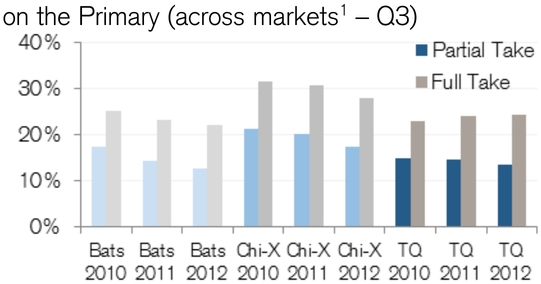

Source: Credit Suisse AES Analysis, Q3 2010, Q3 2011 and Q3 2012

What does it mean for trading?

One way to minimise price fade – supported by the data – is to use "smart take" functionality, which leaves some volume behind rather than taking the whole price point. However, "smart take" also reduces the total amount of liquidity available per trade. As such, this strategy trades immediacy of liquidity for potential reduction in fade.

In some scenarios (e.g. when trying to trade oversized orders quickly), it may make sense to send aggressive orders to each venue straight at the top limit, rather than ticking up the book. This "limit sweep" tactic minimises the chance of price fade on each venue by denying counterparties at deeper prices the opportunity to cancel their orders; however, it does not take full advantage of iceberg liquidity. More details on AES's "Limit Sweep" functionality as well as "smart take" are provided from page 7.

Venue Fade: Cancellations on a Different Venue

Fade can also spread across venues, with trading on Venue A leading to cancellations on Venue B. This could be caused by high frequency traders reacting quickly to cancel orders on other venues before any trades can occur there. Exhibit 15 shows a real example of this on Unilever where:

1) Buying the entire volume available on Chi-X at 2284 leads to
2) 384 shares being cancelled on Bats, and
3) 4742 shares cancelled on the LSE before any further trades occur.

Exhibit 16 demonstrates that a "partial take" on the primary leads to volume disappearing on Chi-X ~17% of the time, with a full take on the primary leading to volume disappearing on Chi-X ~28% of the time (further breakdowns are provided in Appendix 3). As with price fade, we find that venue fade is less likely during times of the day when spreads are higher, and that the probability is marginally increased after the US open.

However, venue fade is not necessarily malicious - smart order routers may redistribute posted volume to a venue that sees an execution, knowing that in the algo world lightning often *does* strike twice. To facilitate this, volume on other venues will be cancelled. The key difference in this scenario is that the "disappearing" volume returns – albeit on a different venue.

Venue Fade – slightly less likely vs 2011, 2010

As with price fade, over the last two years full takes have produced higher probabilities of venue fade than partial takes (see Exhibit 17). However, the overall likelihood of seeing venue fade after a primary trade has decreased slightly, with the exception for fade on Turquoise following "full takes".

This contrasts with the increasing likelihood of price fade shown in Exhibit 14. It may be that increased speed of players on the same venue (e.g. colocation) is driving the increase in price fade, while the SOR redistributions often behind venue fade have stayed relatively similar or become more optimised. HFT players may also now be more active on primary markets – which anecdotal evidence suggests – reducing the likelihood of collateral fade on MTFs.

Adapting to Venue Fade

Many of the same trade-offs found with price fade (e.g. price point preservation vs immediacy) apply to venue fade. However, AES has developed "Blast" to deal with the specific challenges presented by coordinating between venues. "Blast" minimises other traders' ability to cancel their orders between your trades on multiple venues. It can be combined with "Limit Sweep" for those seeking the most aggressive takes (further details are provided from page 7).

What is Momentum Ignition?

Momentum ignition refers to a strategy that attempts to trigger a number of other participants to trade quickly and cause a rapid price move.

Why Trigger Momentum Ignition?

By trying to instigate other participants to buy or sell quickly, the instigator of momentum ignition can profit either having taken a pre-position or by laddering the book, knowing the price is likely to revert after the initial rapid price move, and trading out afterwards.

Momentum Ignition

Likelihood and Rapid Price Moves

Momentum ignition does not occur in the blink of an eye, but its perpetrators benefit from an ultra-fast reaction time. Generally, the instigator takes a pre-position; instigates other market participants to trade aggressively in response, causing a price move; then trades out. We identify momentum ignition with a combination of factors, targeting volume spikes and outsized price moves - see Exhibit 18 for a example of this pattern in Daimler on 13th July, 2012:

Exhibit 18: Daimler AG, 13th July, 2012

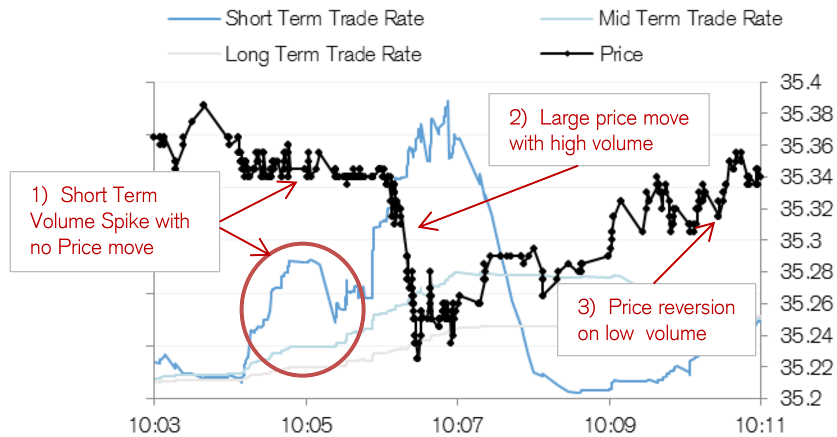

Source: Credit Suisse AES Analysis

To pinpoint momentum ignition, we search for:

1) Stable prices and a spike in volume (Box 1 in Exhibit 23)
2) A large price move compared to the intraday volatility (Box 2)
3) Reversion (Box 3)

Though we cannot conclusively determine the intention behind every trade, this is the kind of pattern we would expect to emerge from momentum ignition. We use this as a proxy to estimate the likelihood and frequency of these events (further details are provided in Appendix 4).

Exhibit 19: Frequency of Momentum Ignition Patterns

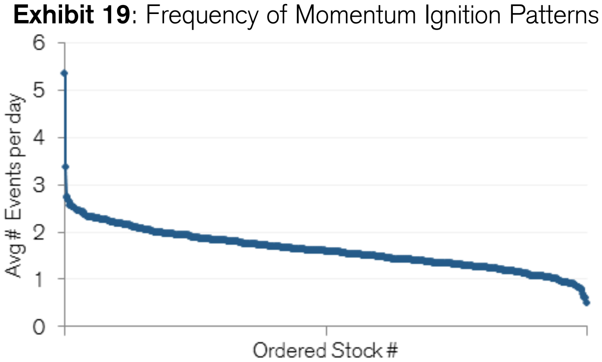

Source: Credit Suisse AES Analysis, STOXX 600 Jul - Sep 2012

Exhibit 20: Price Move Distribution (Post Momentum Ignition Pattern)

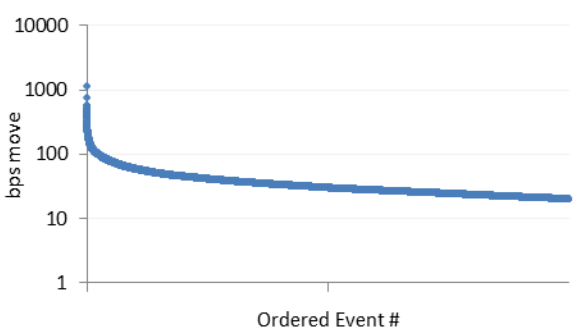

Source: Credit Suisse AES Analysis, STOXX 600 Jul - Sep 2012

Exhibit 21: Duration Distribution (Momentum Ignition Patterns)

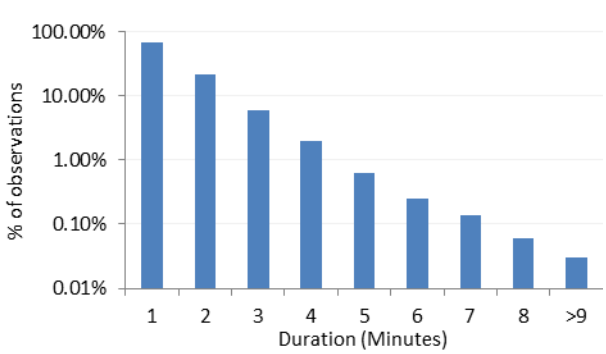

Source: Credit Suisse AES Analysis, STOXX 600 Jul - Sep 2012

Likelihood and Rapid Price Moves

As shown in Figure 19, we estimate these patterns occured on average 1.6 times per stock per day for STOXX 600 names in Q3 2012, with almost every stock in the STOXX600 exhibiting this pattern on average once a day or more. In addition, we note that the average price move is 38bps (but over 5% are more than 75bps, with some significantly higher – see Exhibit 20), and the time it takes for that move to occur is approximately 1.5 minutes (see Exhibit 21). While 38bps may not sound like a big move, it is a bit more significant when compared to the average duration of these events (1.5 minutes) and the average spread on the STOXX600 (approximately 8bps).

Though not all momentum ignition events result in massive price moves, those that do can cause significant impact. Percentage of volume orders that would normally execute over hours may complete in minutes on the back of "false" volume (one of the causes of the 2010 flash crash was a straightforward percentage of volume order). AES offers a variety of protections to help mitigate this kind of dislocation, including customised circuit breakers, active limits (that kick in when the stock decouples from a specified index) and fair value limits (more details in the next section).

Exhibit 22: Quote Filtering: Heineken, 2nd May, 2011

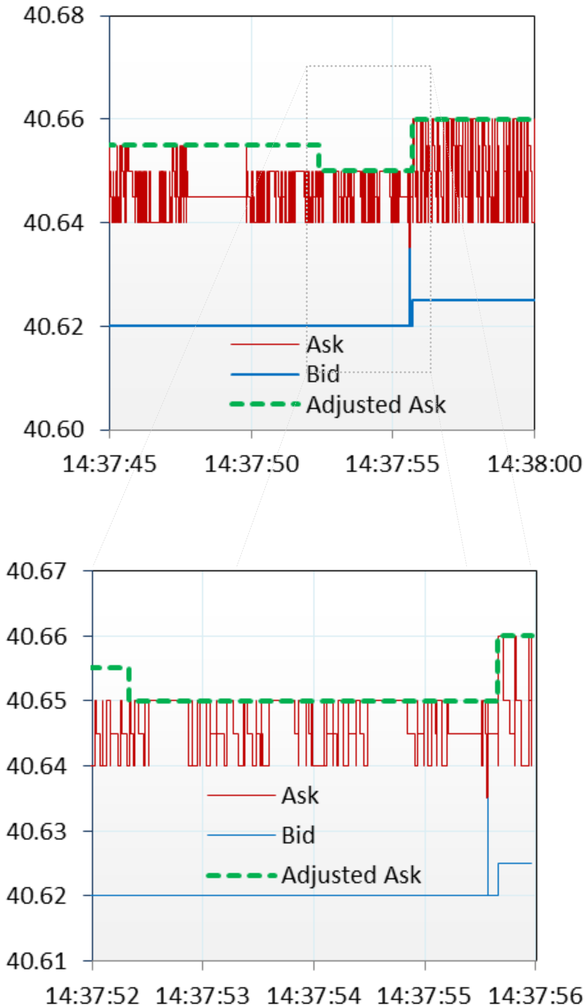

Source: Credit Suisse AES Analysis

What Does AES do about it?

Adaptive Behaviour and Quote Filtering

As mentioned in the previous section, a variety of techniques – including pattern recognition, burst detection and feature extraction – are used to detect various negative HFT behaviours and adapt accordingly.

For instance, we have designed quote filtering methodology to score (and subsequently flag) potential HFT activity, and generate an 'adjusted bid' and 'adjusted ask' (i.e. with the HFT quote stuffing removed), rather than only the 'HFT affected' quotes (see Exhibit 22).

It is also designed to distinguish updates due to "real" market trades from excessive updates purely generated by HFT activity. This way, rapid price (and quote) moves – driven by trading on the back of news, for example – are not mistakenly flagged as quote stuffing scenarios. In the absence of such "real" trading, the quote filtering logic would kick in.

Whilst passive orders might prefer adjusted quotes, aggressive strategies may attempt to take advantage of what appears to be fleeting liquidity. If the opposite side of the spread is coming towards the order, it may try to "pick off" that transient volume.

To do this, aggressive strategies will only ever send IOCs and only when the temporarily 'narrowed' price is one which the strategy would wish to trade at. Strategies are never induced to pay worse prices by the existence of fleeting quote activity.

Enhanced Functionality in Guerrilla (and other tactics)

AES's Guerrilla tactic takes certain HFT activity into account when determining fair value levels, aggressiveness and trading behaviour. It may then adjust various parameters to alter its behaviour to enhance the intelligence of its trading, including 'pausing' the order in certain situations. Other tactics also have access to this logic, with each tactic adjusting its behaviour in a unique way.

Dark-only flow traded through AES (e.g. in tactics such as Crossfinder+) can minimise the chance of being affected by 'mid-point gaming' with by withdrawing from certain venues, raising MAQs and using tighter limits. These protections will allow the midpoint to come towards the order – enabling the strategy to participate at a temporarily more favourable price – but restrict it from moving away.

If apparent gaming occurs consistently on a particular venue or with a particular counterparty in Crossfinder, the AES Alpha Scorecard will pick this up and highlight that venue or that counterparty as exhibiting excessive "opportunistic" behaviour[1]. Credit Suisse's clients then have the ability to decide whether to trade on those venues or against that group of counterparties. See Classifying Dark Counterparties for more details on the AES Alpha Scorecard and how we use it to quantify dark pool trading.

[1] However, a counterparty being opportunistic does not necessarily imply they have been attempting to take advantage of mid-point gaming, as other (perfectly valid) trading styles can also lead to such a classification.

Table 1: Credit Suisse AES European SOR Configuration Options

	Tick Sweep (Walk up the book when aggressing on each venue)	**Limit Sweep** (Go to each venue at the ultimate limit)
Blast (Send specifically targeting venue fade)	+ Minimise venue fade, take advantage of iceberg liquidity - Slower than Parallel Tick Sweep, can miss prices in fast moving markets.	+ Target price and venue fade at all costs - Slower than Parallel Limit Sweep. Can trade at worse prices by missing hidden liquidity
Parallel (Send to each venue in parallel)	+ Trade quickly, taking advantage of hidden liquidity - Not as direct as Parallel Limit Sweep meaning slight risk of fade	+ Fastest to market, minimise price fade - Does not take full advantage of icebergs - may trade at worse prices as a result
Serial (Visit each venue in turn)	+ Check every price point on every venue for hidden liquidity - Slowest option and at risk of price and venue fade	+ Minimise price fade on each venue in turn - High risk of taking worse prices than available on other venues

Flow that reaches Credit Suisse's dark pool (Crossfinder) via aggregators does not receive such protections, as Crossfinder is simply an execution venue for this flow. When interacting through AES algorithms, these additional protections are available.

Targeting Fade with AES Blast and Limit Sweep

Order Book Fade is clearly present in Europe, and one might suppose that configuring an execution strategy to minimise fade would be the best way to extract liquidity. However, over-emphasising fade can – in some scenarios – result in sub-optimal executions. AES, in conjunction with our SOR, can be customised to balance the trade-off between targeting fade and intelligently searching for extra liquidity.

Using AES "Limit Sweep" – which sends to all venues immediately at the top limit – helps combat price fade, but it can mean not taking full advantage of icebergs. AES has also introduced "Blast" functionality to specifically target venue fade, mitigating the advantage that HFT firms try to take exploit; however, this approach is at the cost of ultimate speed.

Rather than dictate to clients exactly which configuration best suits their goals, Credit Suisse provides the ability for any of the options in Table 1 to be configured on a tactic by tactic basis[2] (as well as for DMA). Each represents a different mix of trade-off between speed, avoiding fade and uncovering hidden liquidity. Both Blast and Limit Sweep have been available since 2011, and a wide variety of combinations have been set up for clients.

Using Smart Take to Avoid Destroying Price Points

Smart Take functionality is available for AES's aggressive strategies: Guerrilla and Sniper. It leaves some volume on the order book – rather than taking out the entire price point – to minimise the price signal being sent to other participants. As demonstrated above, this reduces the occurrence of fade.

Smart Take is currently offered as a customisation rather than a default. This is due to the trade-off between reduced fade and immediacy of liquidity mentioned earlier. For those who want to make use of the functionality, it can easily be configured.

Dynamic Fair Value, Active Limits and Custom Circuit Breakers

While momentum ignition strategies may not send a "high frequency" signal, AES nevertheless protects our clients from any significant price moves by providing a Dynamic Fair Value protection on all orders - including dark only orders. On by default for all relevant headline strategies[3], Dynamic Fair Value also helps protect AES clients from other scenarios, such as fat fingers.

Similarly, during momentum ignition a stock may temporarily decouple from a related index. AES Active Limits can be configured to restrict trading in this scenario, preventing the intended gaming. Additionally, AES allows for custom circuit breakers, which automatically pause any orders in stocks that move more than a certain percentage away from their arrival price. See Enhancing Protections and Transparency in Europe for further details on these options.

[2] For example, by default, "Limit Sweep" is used on Sniper Aggressive, with Normal and Patient Sniper using tick sweep.
[3] i.e. Excluding Sniper and Reserve

Conclusion

While our previous report (High Frequency Trading – The Good, The Bad, and The Regulation) noted that not all HFT strategies are negative, it is clear that certain undesirable HFT behaviour does exist. In this piece, we have presented a number of concrete examples of "bad" HFT – including quote stuffing, order book fade and momentum ignition – as well as some broader details on the prevalence of those HFT strategies.

We found, for instance, that quote stuffing occurs much more frequently on MTFs. This could mean that dark pools referencing the EBBO are more susceptible to mid-point gaming. Additionally, we saw that likelihood of price fade has increased since 2010 (though venue fade has decreased slightly). Finally, we showed that momentum ignition – though hard to isolate – can trigger outsized price moves in short periods of time.

AES employs a variety of signal processing techniques to detect and categorise HFT behaviour in the market (which we have also used to analyse the data presented in this report). Armed with this real time information, we dynamically adjust the behaviour of our algorithms to maximise the protection afforded and take advantage of any favourable scenarios that may result. For example, opportunistic tactics such as Guerrilla and Sniper have been continually enhanced to reflect the changing trading environment. We have also developed additional functionality – such as Quote Filtering, Blast and Limit Sweep – which can also be integrated with other strategies.

With the influence of High Frequency Trading likely form a significant part of the trading landscape for the foreseeable future, it is paramount to remain aware and informed about HFT. Rather than branding all HFT as good or bad, we should strive to adapt: protecting against the negative behaviours while we take advantage of those that add value.

Appendix 1: Some Further Example Quote Stuffing Patterns

Our detection methodology has highlighted a variety of patterns in (for example) quote updates when HFT activity has been detected. We present some (zoomed in) examples below.

Square Wave

Exhibit A1: Arkema, 13th July, 2012

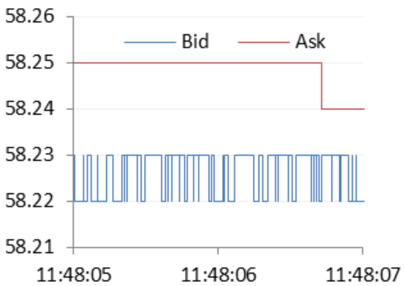

Source: Credit Suisse AES Analysis

Sawtooth

Exhibit A2: Inditex, 3rd October, 2012

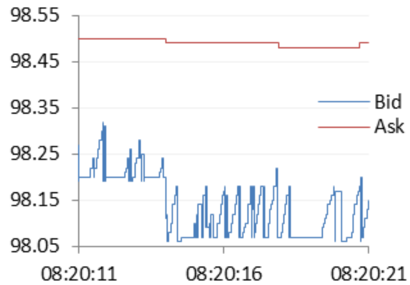

Source: Credit Suisse AES Analysis

Other Patterns:

Exhibit A3: HSBC, 18th July, 2012

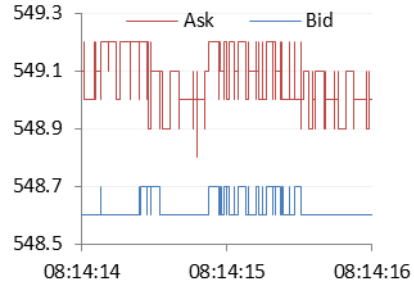

Source: Credit Suisse AES Analysis

Exhibit A4: ASML Holding (Chi-X), 12th July, 2012

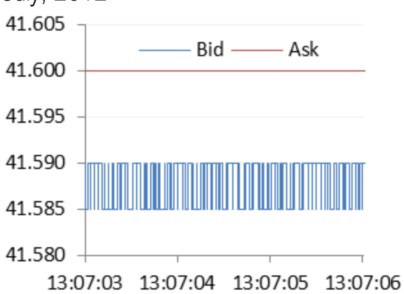

Source: Credit Suisse AES Analysis

Exhibit A5: BBVA, 14th August, 2012

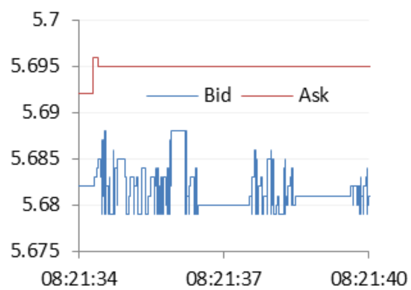

Source: Credit Suisse AES Analysis

Exhibit A6: Veolia Environnement, 18th June, 2012

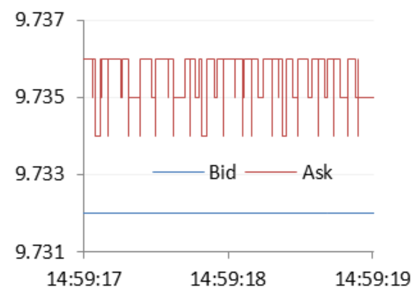

Source: Credit Suisse AES Analysis

Appendix 2: HFT Detection and Response

To outline one of the methods by which we detect certain HFT behaviour in real time, we define a function f which generates a 2 dimensional HFT measure based on orderbook events. In particular, we (partially) outline our quote stuffing detection and filtering logic below: Specifically, define $f(\boldsymbol{q_a}, \boldsymbol{q_b}, \boldsymbol{\tau}, \boldsymbol{q}^*, \tau^*, \ldots, s, t)$ such that

$$f: A \rightarrow [0, \infty) \times [0, \infty)$$

is a function that takes real time market data combined with historic baseline numbers and returns a 2-D measure, where
A represents the space of possible input vectors (i.e. the possible values of $\boldsymbol{q_a}, \boldsymbol{q_b}, \boldsymbol{\tau}, \boldsymbol{q}^*, \tau^*$ etc.),

s, t are time parameters,

$\boldsymbol{q_a}, \boldsymbol{q_b}$ are orderbook update events on the ask side and bid side of the book respectively, potentially across venues

$\boldsymbol{\tau}$ represents trade updates

$\boldsymbol{q}^*, \tau^*$ are 'baseline' values for q_a (and q_b) and τ respectively, and

$[0, \infty) \times [0, \infty)$ represents the possible HFT measure values for the ask and the bid.

Whilst we do not explicitly define f in this article - nor all of its input parameters (in part to avoid it being reverse engineered and 'subverted'), we use various feature extraction techniques (including burst detection methods) to obtain our results.

We also define stock specific thresholds f^* such that if $\|f\|_\infty > f^*$ then $HFT = 1$, and further, if $\|f(\boldsymbol{q_a}, \boldsymbol{0}, \ldots)\|_\infty > f^*$ then $HFT_{ask} = 1$ and if $\|f(\boldsymbol{0}, \boldsymbol{q_b}, \ldots)\|_\infty > f^*$ then $HFT_{bid} = 1$.

As demonstrated in the various charts, we again use techniques adapted from signal processing to create a stable ask and stable bid where $g_{ask}: A \rightarrow [0, \infty)$ and $g_{bid}: A \rightarrow [0, \infty)$ where g is a function of the ask (or bid) price movements as well as a function of f, and returns the stable ask or bid as highlighted in the green dotted lines in e.g. Exhibit 6. Again, we will not explicitly define g in this article, but note that various pattern recognition techniques are employed to enhance our results.

* More accurately, g_{ask} and g_{bid} map from A to a (countably infinite) subset of $[0, \infty)$ i.e. onto the permitted prices as defined by tick size limitations, which vary by stock and market.

Appendix 3: Measuring Fade and Asynchronous Market Data

When analysing the effect and prevalence of fade in the market, we used our tick database to combine quote and trade data from multiple venues to isolate and estimate the prevalence of both price and venue fade. This requires synchronisation across price feeds and subsequent classification of any fade detected – we used a variety of 'pattern matching' rules to do this.

These rules are a multi-dimensional function of the various tick data streams and are both forward and backwards looking in time (depending on market). For some markets (on which the data presented in the main body of the text is based), the sequencing of quote and trade updates are sufficiently in sync that we are able to determine the presence or absence of fade (and the type if any) – and we present some further breakdowns here. Price fade (i.e. volume fade on the same venue the trade occurred) is represented by the lighter coloured bars, with venue fade also shown (fade on a venue different from traded venue). Interestingly, there does appear to be a difference between the Scandinavian markets compared to the LSE for example, which could well be a consequence of microstructural differences such as tick size and average queue length making fade less likely on Scandinavian Markets.

Table A1: Extended Order Book Fade Data

Across Markets	**LSE**	**Scandi (aggregated)**

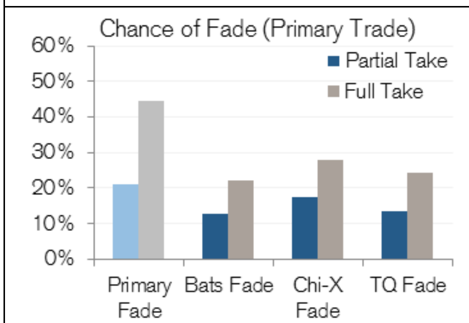

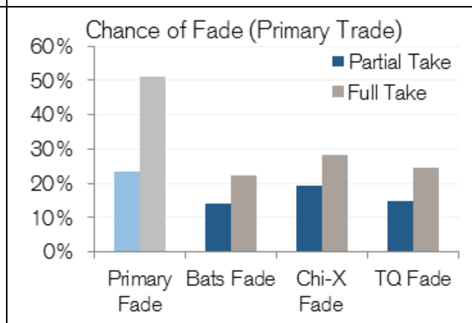

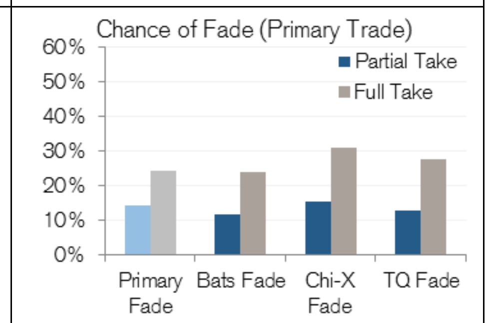

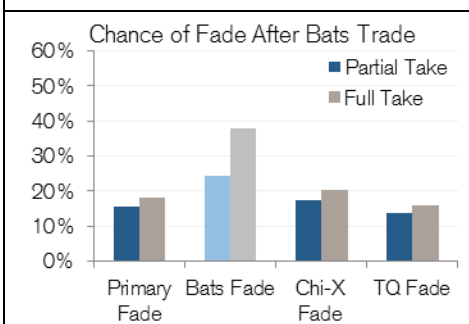

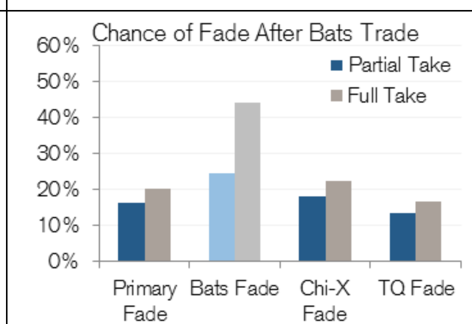

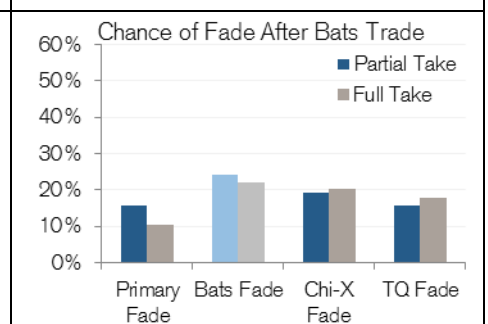

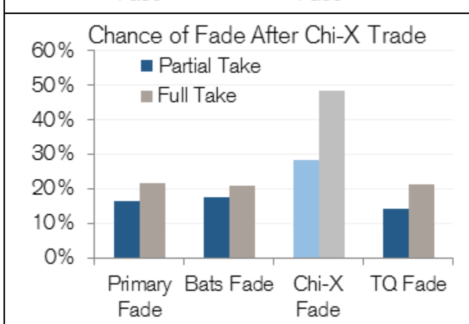

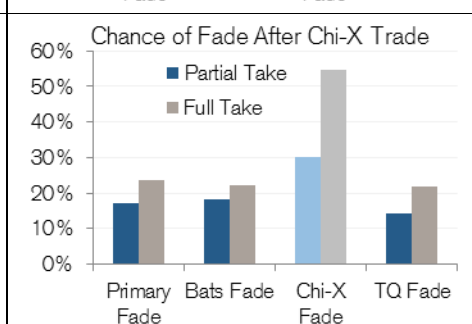

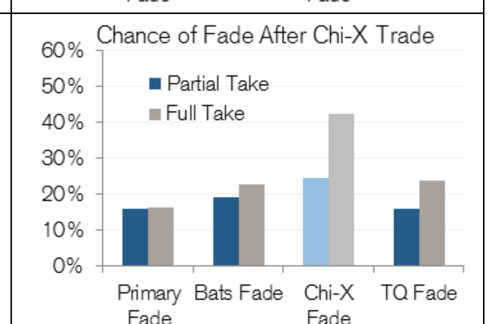

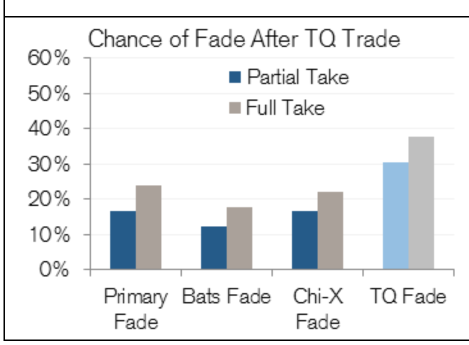

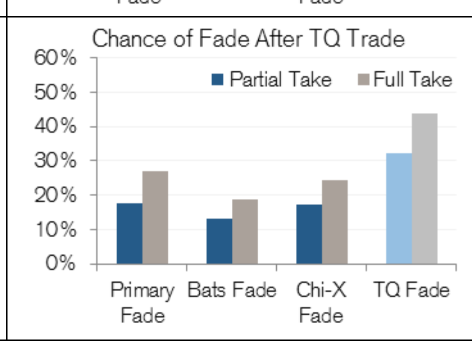

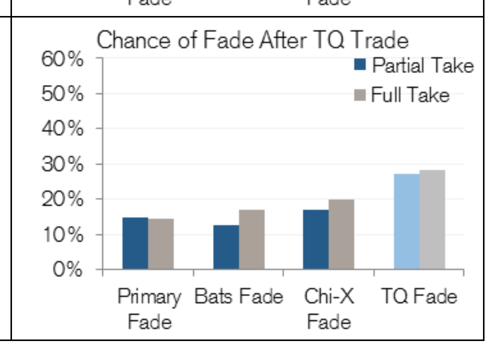

Source (all charts on page): Credit Suisse AES Analysis, Jul – Sep 2012
Across Markets = Copenhagen, Helsinki, London, Madrid, Milan, Oslo and Stockholm. Scandi (aggregated) = Copenhagen, Helsinki, and Stockholm.

Appendix 3: Measuring Fade and Asynchronous Market Data (contd..)

However, data from other markets are not so clean, and can mean that it is essentially impossible to reliably uncover the exact sequencing of events in any programmatic fashion, meaning that the pattern matching rules may return correct results in some scenarios but spurious results in others.

Exhibit A7 provides an example of quote and trade updates being well out of sync, with the trade record (900 shares @59.75) occurring well after the quote update showing 900 bid at 59.75 disappearing, with multiple other quote updates in between. This makes it very difficult to isolate the actual order of events and hence the historic presence or absence of fade efficiently.

While in this scenario it is straightforward to determine which quote update corresponds to the trade update, this is just one of the simpler examples of such asynchronous updates meaning certain markets (e.g. Euronext, Swiss, Germany) are excluded from the aggregated analysis provided, as too many spurious results are returned without extra manual intervention and validation on historic analysis.

Nonetheless, Blast, Limit Sweep and Smart Take are still available as options on these markets as the asynchronous nature of the market data does not invalidate the benefits these protections can provide (asynchronous data merely affects the historic detection of these events).

Exhibit A7: Example of Asynchronous Updates Nestle (Virt-X), 10th October, 2012

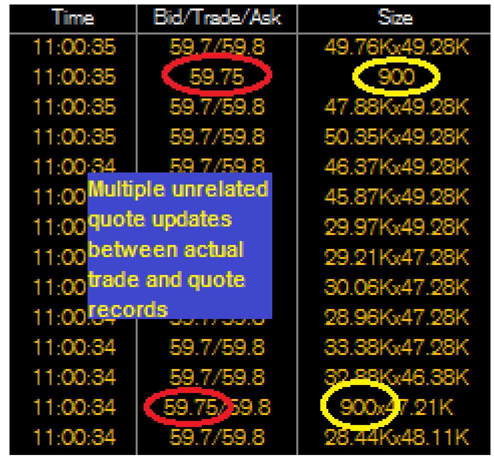

Time	Bid/Trade/Ask	Size
11:00:35	59.7/59.8	49.76Kx49.28K
11:00:35	59.75	900
11:00:35	59.7/59.8	47.88Kx49.28K
11:00:35	59.7/59.8	50.35Kx49.28K
11:00:34	59.7/59.8	46.37Kx49.28K
11:00[illegible]	[illegible]	45.87Kx49.28K
11:00[illegible]	[illegible]	29.97Kx49.28K
11:00[illegible]	[illegible]	29.21Kx47.28K
11:00[illegible]	[illegible]	30.06Kx47.28K
11:00[illegible]	[illegible]	28.96Kx47.28K
11:00:34	59.7/59.8	33.38Kx47.28K
11:00:34	59.7/59.8	32.88Kx46.38K
11:00:34	59.75/59.8	900x47.21K
11:00:34	59.7/59.8	28.44Kx48.11K

Source: Credit Suisse AES Analysis

Appendix 4: Detecting Momentum Ignition Patterns

When using our proxy detection logic for momentum ignition, as stated in the text, we essentially look for a spike in volume accompanied by no price move, followed by a significant price move and then reversion or stability. In more technical terms, we employ a multistage filter to identify these events:

Denote a momentum ignition style pattern starting at time t if:

$$f(\boldsymbol{v}(t-s_1,t)) > a * \boldsymbol{v}(t-s_2,t) \text{ and } g(\boldsymbol{v}(t-s_1,t)) > b * \boldsymbol{v}(t-s_3,t) \text{ and}$$

$$abs\big(p(t-s_4)-p(t)\big) < c * \sigma_{intraday} \text{ and } abs\big(p(t+s_4)-p(t)\big) > d * \sigma_{intraday}$$

where f, g are functions that provide adjustment factors,

$\boldsymbol{v}(u,v)$ denotes the average trade rate over the period (u,v),

t represents a point in time,

s_n represent various timeframes (with $s_1 < s_2 < s_3$),

a, b, c, d are constants and

$\sigma_{intraday}$ represents intraday volatility.

Then denote the end of the ignition period t^* (but not the end of the reversion period), and define it such that

$t^* = \min\{\, u \mid u > t + s_4$ and $\mathrm{sgn}(p(t^* + s_4) - p(t)) = -\mathrm{sgn}(p(u) - p(u - s_4))\}$

Market Commentary Disclaimer