HIGH FREQUENCY TRADING

Table of Contents

CHAPTER 1: UNDERSTANDING TODAY'S FUTURES LIMIT ORDER BOOK

What is the Limit Order Book?

The cumulative trade size of all available limit orders that are able to meet incoming market orders on a particular trading venue at any given time is known as ***liquidity.*** Having a larger number of limit order traders available on the exchange, with each trader having a larger size of limit orders, means more liquidity for the given trading venue. Liquidity is also inevitably limited in today's markets – the number of limit orders is quantifiable, and each limit order has a definite size (Aldridge 2013).

The majority of contemporary exchanges are structured as so-called centralized limit order books (CLOBs), also known as a double-sided auction, to account for limit orders. The CLOBs were established in the United States in the early 1970s and implemented in Europe in the 1980s. In a CLOB model, all incoming limit orders are documented in a "book" which contains a table with columns that correspond to sequential price increments and rows that reflect the sizes of limit orders posted at every price increment. The information in this limit order book can be distributed to all other market participants as Level II data or Depth of Market Data.

In theory, limit order books are usually assumed to be symmetrical with the market price, with the distribution of limit buy orders mirroring that of limit sell orders. Moreover, in several risk management applications, order books are also assumed to follow normal bell-curve distributions. However, neither of these assumptions tends to hold, as order books are rarely normal and are often asymmetric.

Market making is a set of high frequency trading strategies involving the placement of a limit order to sell (or offer) or a buy limit order (or bid) for the purpose of earning the bid-ask spread. With these strategies, market makers deliver counterpart to incoming market orders. The role of market maker was traditionally performed by specialist firms, but this class of strategy is now applied by a wide range of investors, mainly due to the wide adoption of direct market access. As pointed out by empirical studies, this renewed competition among liquidity providers has led to the reduction of effective market spreads, and, consequently, the reduction of indirect costs for final investors (Hendershott, Jones, and Menkveldf 2011).

Figure 1.1 - View of a Typical "Centralized Limit Order Book" (CLOB).

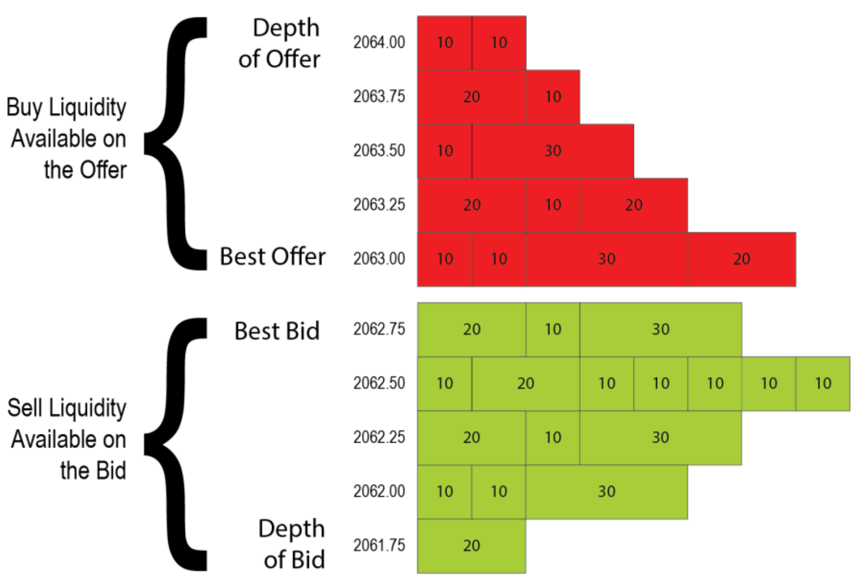

Figure 1.2 shows that when a new limit order arrives, it is placed into a limit order queue corresponding to its price. Since all the prices in today's markets are subject to a minimum increment, or tick, price-based bins are clearly defined. The limit buy orders at the highest price form the best bid, with the price of these orders reflected in the best bid price, and the aggregate size reported as the best bid size. Correspondingly, the limit sell orders posted at the lowest price form the best ask, with their respective price and size information. Best ask is sometimes referred to as best offer. At any given moment of time, there exists a fixed aggregate size of all limit orders posted at each price.

Figure 1.2 –The order structure in a “First In, First Out” (FIFO) market. See facing page.


2064.00 10 10
2063.75 20 10
2063.50 10 30
2063.25 20 10 20
Best Offer 2063.00 10 10 30 20
Best Bid 2062.75 20 10 30
2062.50 10 20 10 10 10 10 10
2062.25 20 10 30
2062.00 10 10 30
2061.75 20
100
New Market Sell Order
for 100 hits the Bid

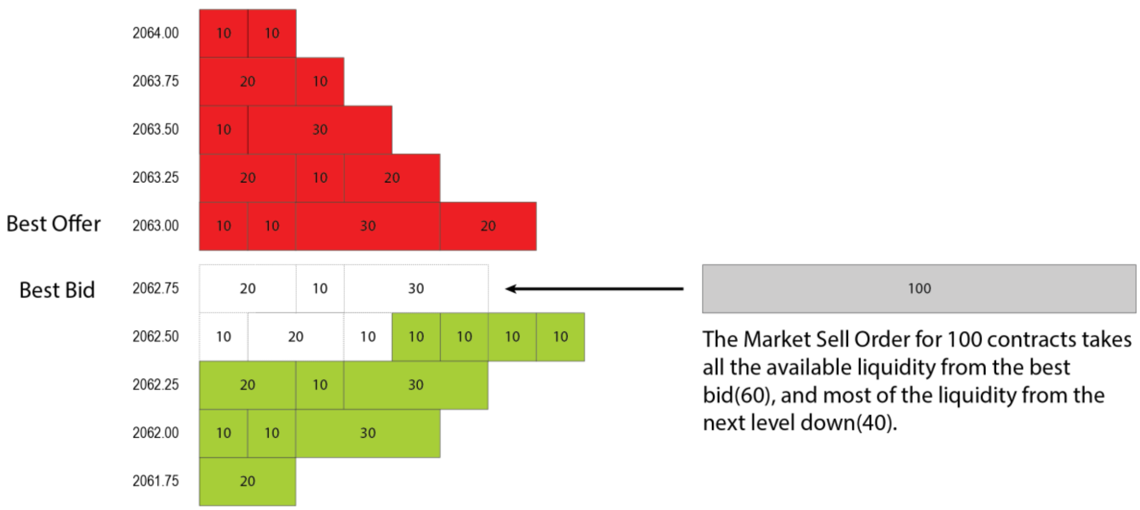
2064.00 10 10
2063.75 20 10
2063.50 10 30
2063.25 20 10 20
Best Offer 2063.00 10 10 30 20
Best Bid 2062.75 20 10 30
2062.50 10 20 10 10 10 10 10
2062.25 20 10 30
2062.00 10 10 30
2061.75 20
100
The Market Sell Order for 100 contracts takes all the available liquidity from the best bid(60), and most of the liquidity from the next level down(40).

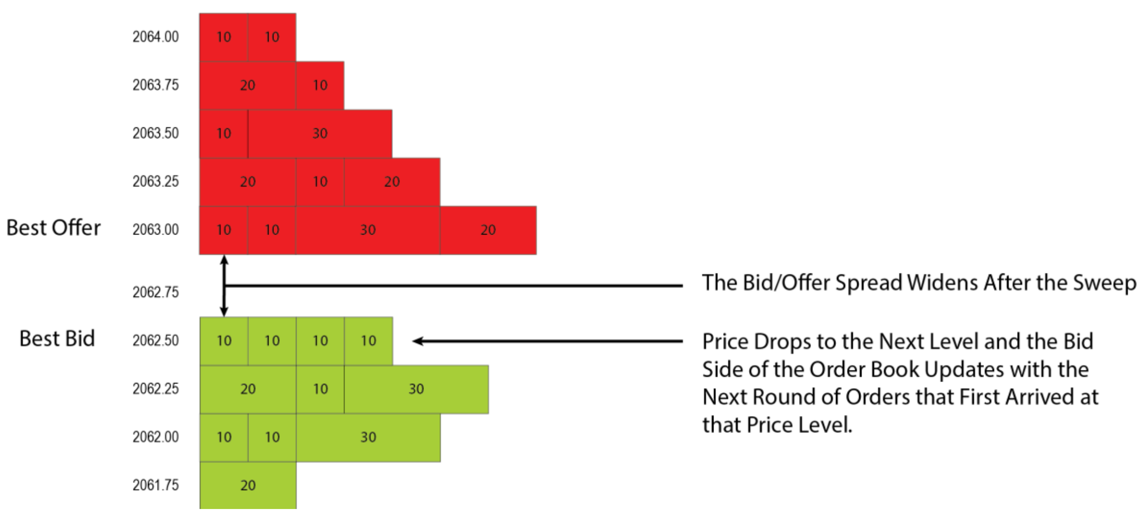
2064.00 10 10
2063.75 20 10
2063.50 10 30
2063.25 20 10 20
Best Offer 2063.00 10 10 30 20
2062.75
The Bid/Offer Spread Widens After the Sweep
Best Bid 2062.50 10 10 10 10
2062.25 20 10 30
2062.00 10 10 30
2061.75 20
Price Drops to the Next Level and the Bid Side of the Order Book Updates with the Next Round of Orders that First Arrived at that Price Level.

Visualizing the Limit Order Book Data

Generally, price is presented in a price ladder format. See below for a graphic representation explaining the BookMap DOM compared to some of our partners Price Ladder displays.

Figure 1.3 - Orients the TT x Trader MD trader DOM to BookMap.

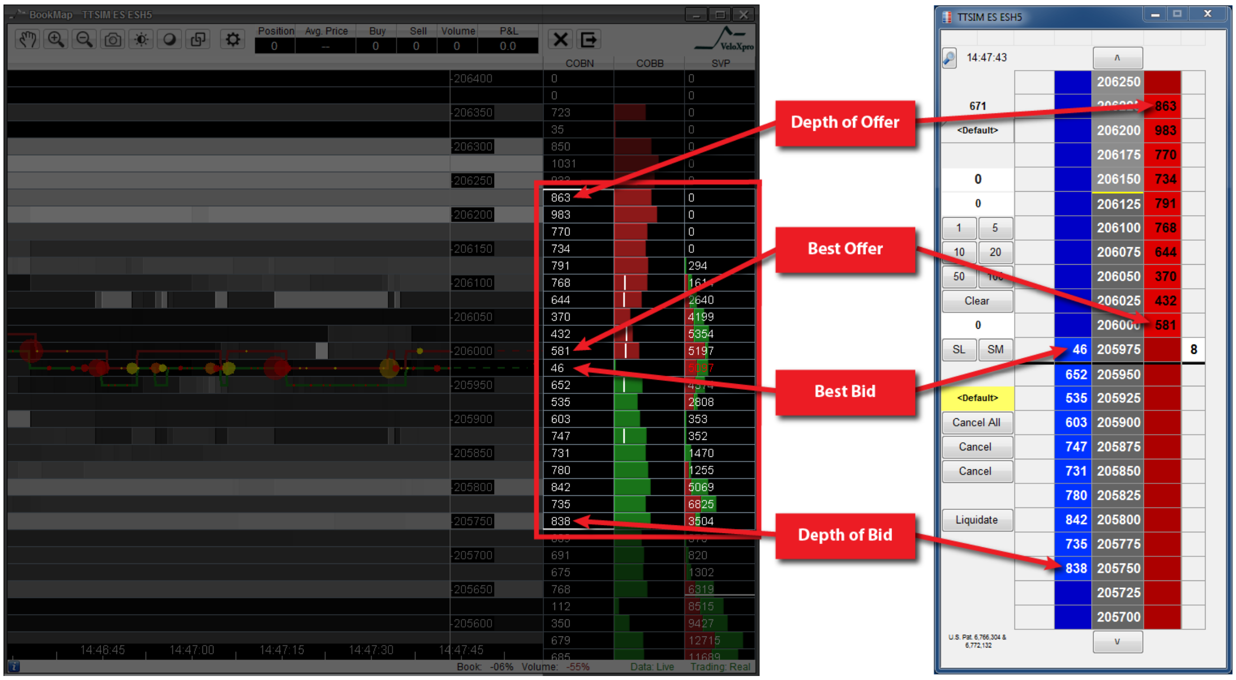

Figure 1.4 - Orients the CQG Trader DOM to BookMap.

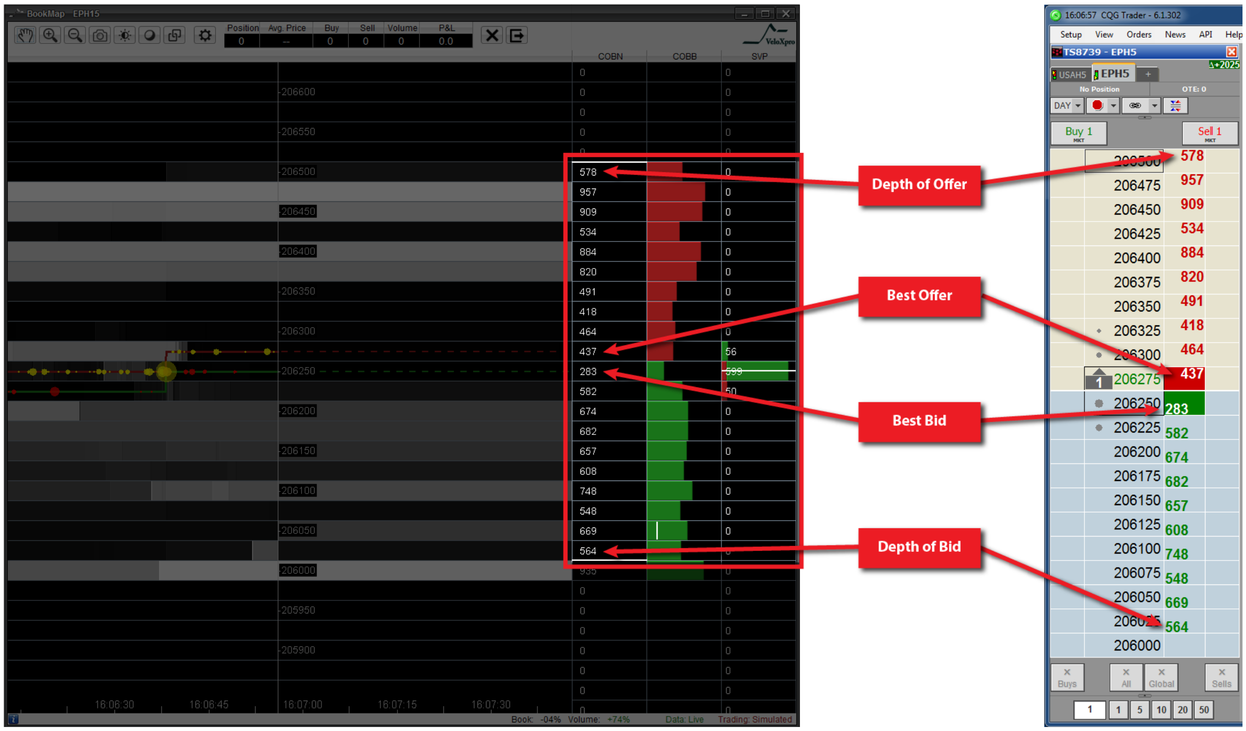

Figure 1.5 - Orients the NinjaTrader DOM to BookMap.

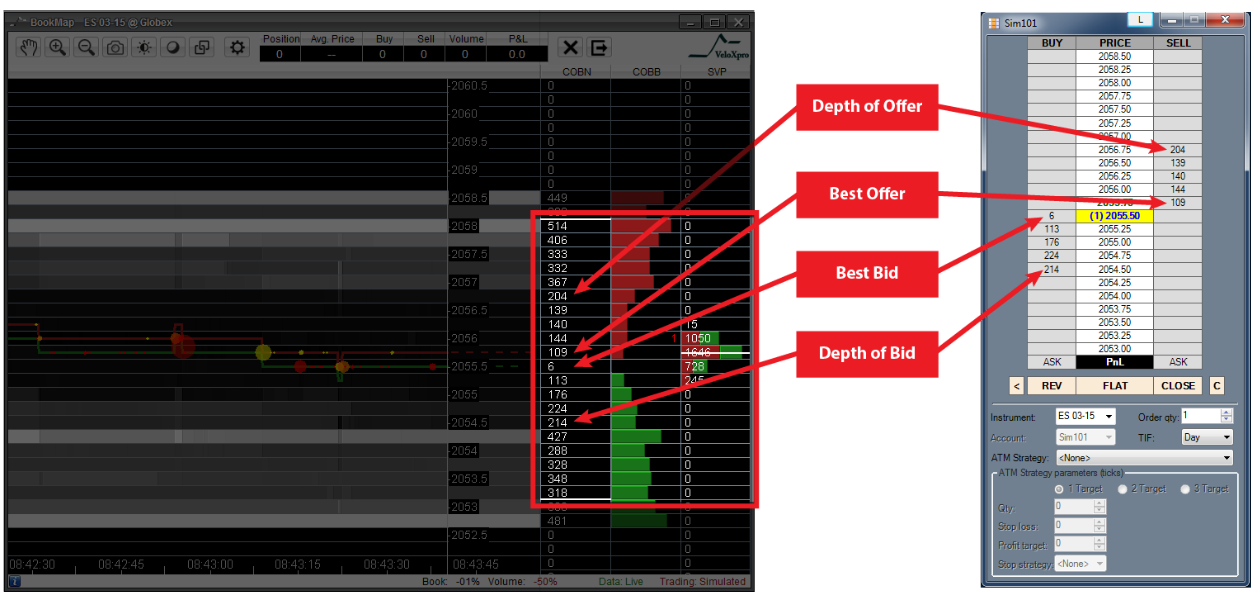

Figure 1.6 - Orients the S5 Trader DOM with BookMap.

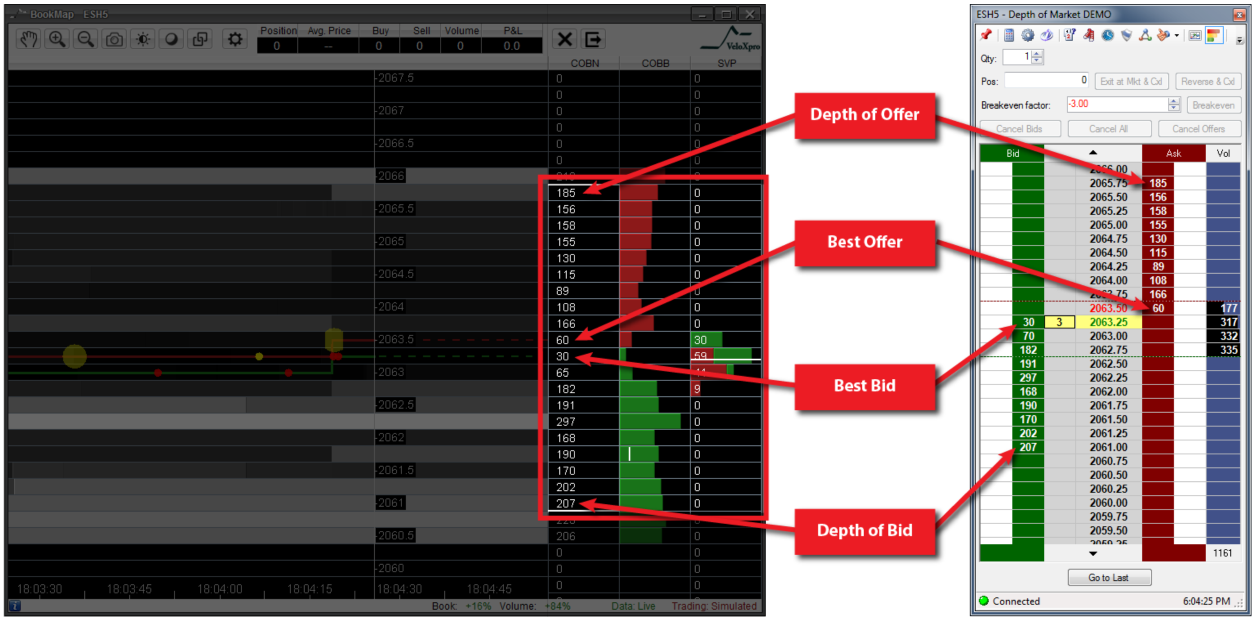

Limit Order Book + HeatMap = BookMap

Normally, the valuable information from the limit order book is lost once a quote refresh or a trade at that level changes the current data. This is where BookMap comes to life. BookMapxRay delivers the only configurable heat map display that shows both real-time order flow and live trade analysis information, combined with historical depth-of-market (DOM) data. With our unique java-based Complex Event Processing (CEP) engine that handles millions of market data updates in real time, it is now possible to see exactly how the limit order book evolves over time. This unique perspective enables traders to get faster and deeper insight into live market dynamics and short-term price action.

In real time, we process ALL Limit Order Book and Trade data. We highlight the variations at the depth using grayscale HeatMap technology, giving you an “X-Ray” type look through the markets and the varied actors in play at any given time.

Figure 1.7 - Illustrates what part of the BookMap is showing Live Data. Note that BookMap utilizes the HeatMap technology in real time on every update from your API.

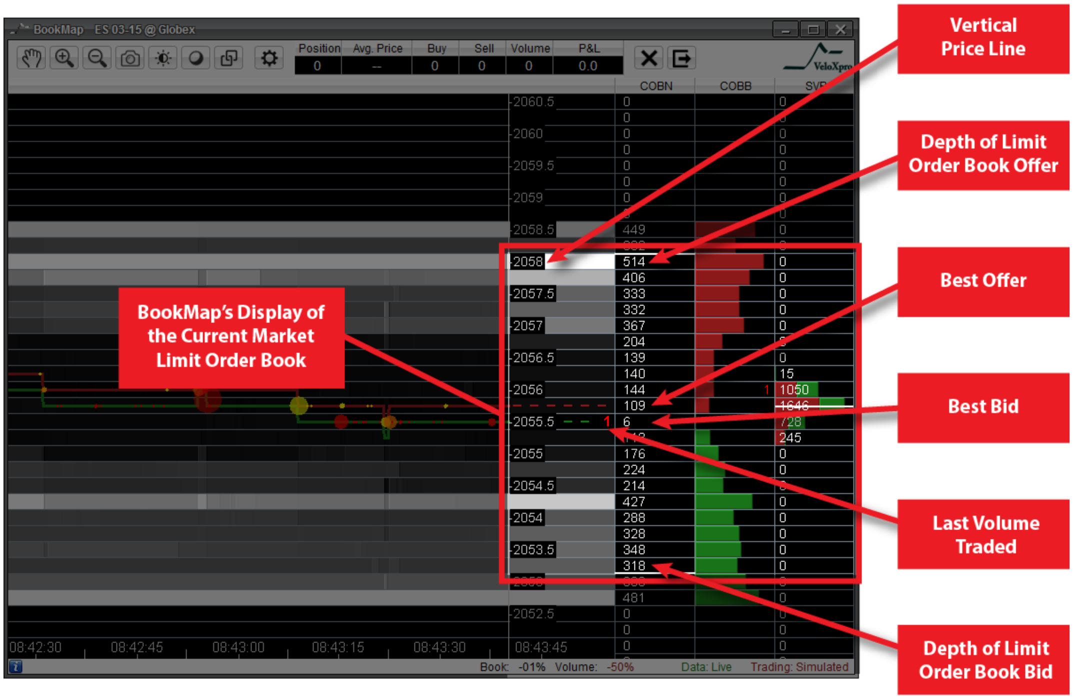

Figure 1.8 - Illustrates the parts of BookMap that display Historical Information.

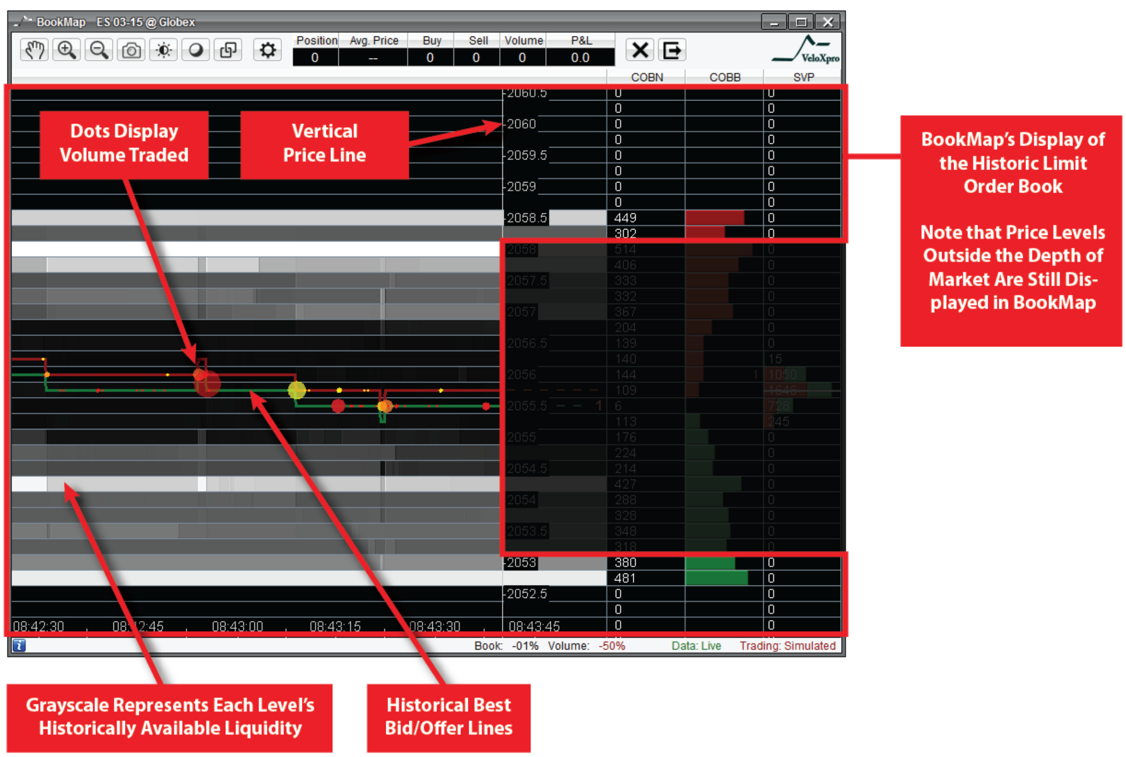

Figure 1.9 - Puts it all together.

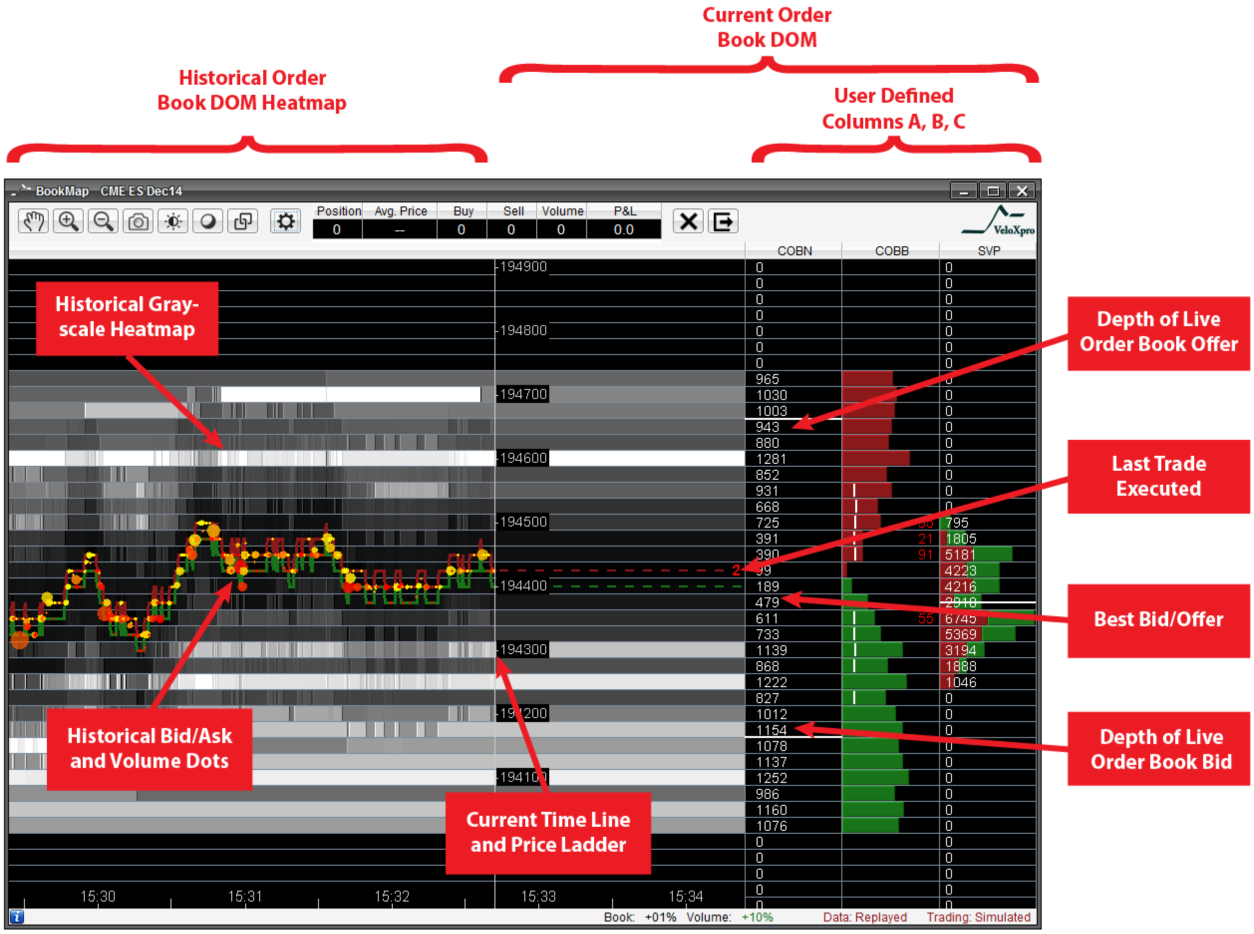

CHAPTER 2: UNDERSTANDING ORDER FLOW

Matching Engines

A market buy order is matched with limit sell orders when it arrives, starting with those placed at the best ask price. If the incoming market buy order's size is greater than the size of the best ask queue, the market order will "sweep" through other offer queues in the direction of increasing price, "eating up" available liquidity at those price ticks. Sweeping creates a significant gap in limit orders on the ask side, instantly increasing the bid-ask spread and possibly inducing slippage in subsequent market buy orders. The order-matching process is similar for market sell orders that end up matched with the available limit buy orders accumulated on the bid size of the book. Limit buy orders with prices equal or higher than the prevailing best bid are executed like market buy orders. In the same way, low-priced limit sell orders are generally treated as market sell orders (Aldridge 2013).

If the incoming buy order's size is smaller than the size of the best ask, and the aggregate best ask queue is composed of numerous limit sell orders placed at the best ask price, the decision as to which of the limit sell orders is matched against the market buy order may vary from one exchange to another. Presently,

most exchanges practice price-time priority, also known as the first-in-first-out (FIFO) execution schedule for limit orders. However, there are several other exchanges that are now implementing pro-rata matching, wherein a fixed proportion of each limit order is matched at a given price.

In time-price priority, or FIFO, execution, the limit order that arrived first is the first of that price bin to be matched with the incoming market order. FIFO is also known as the continuous auction, and has been shown to improve transparency of trading via the following measures:

- Reducing information asymmetry – the limit order book information is accessible to all traders.
- Enhancing liquidity – a CLOB structure motivates traders to add limit orders, which consequently increases market liquidity.
- CLOB's organization supports efficient price determination with its fast and objective order-matching mechanism.
- Standard rules for all market participants guarantee operational fairness and equal access.

Figure 2.1 - Price-Time Priority Execution. See facing page.


2064.00 10 10
2063.75 20 10
2063.50 10 30
2063.25 20 10 20
Best Offer 2063.00 10 10 30 20
Best Bid 2062.75 20 10 30
2062.50 10 20 10 10 10 10 10
2062.25 20 10 30
2062.00 10 10 30
2061.75 20
100
New Market Sell Order
for 100 hits the Bid

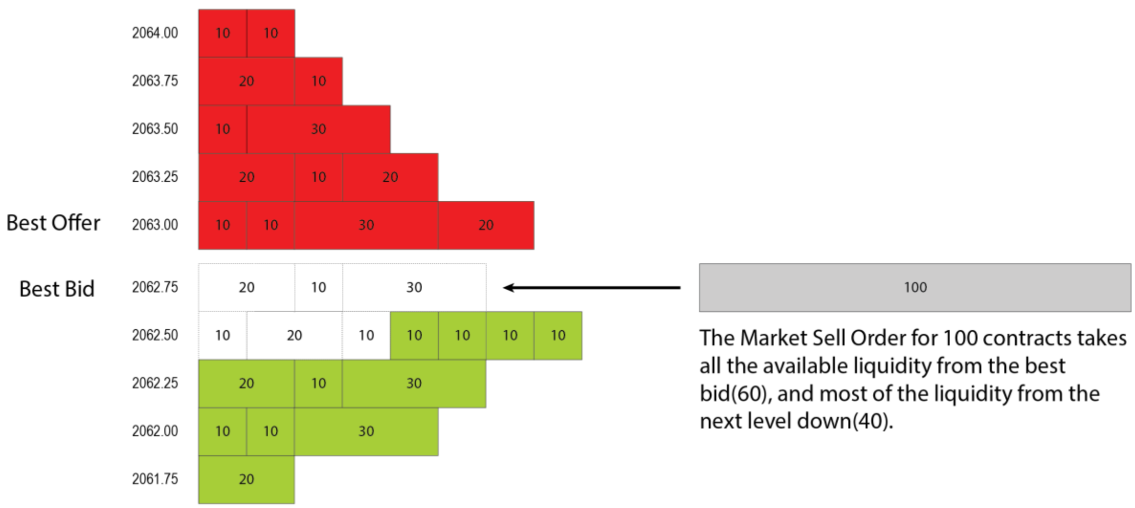
2064.00 10 10
2063.75 20 10
2063.50 10 30
2063.25 20 10 20
Best Offer 2063.00 10 10 30 20
Best Bid 2062.75 20 10 30
2062.50 10 20 10 10 10 10 10
2062.25 20 10 30
2062.00 10 10 30
2061.75 20
100
The Market Sell Order for 100 contracts takes
all the available liquidity from the best
bid(60), and most of the liquidity from the
next level down(40).

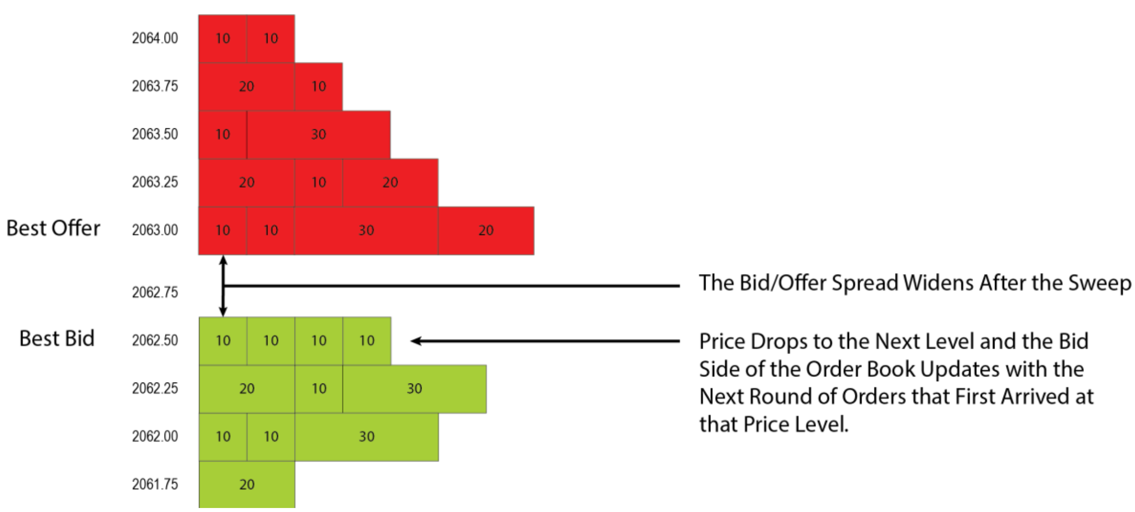
2064.00 10 10
2063.75 20 10
2063.50 10 30
2063.25 20 10 20
Best Offer 2063.00 10 10 30 20
2062.75
The Bid/Offer Spread Widens After the Sweep
Best Bid 2062.50 10 10 10 10
Price Drops to the Next Level and the Bid
Side of the Order Book Updates with the
Next Round of Orders that First Arrived at
that Price Level.
2062.25 20 10 30
2062.00 10 10 30
2061.75 20

Figure 2.2 shows a BookMap representation in real time of a “sweep” through several levels in a FIFO market.

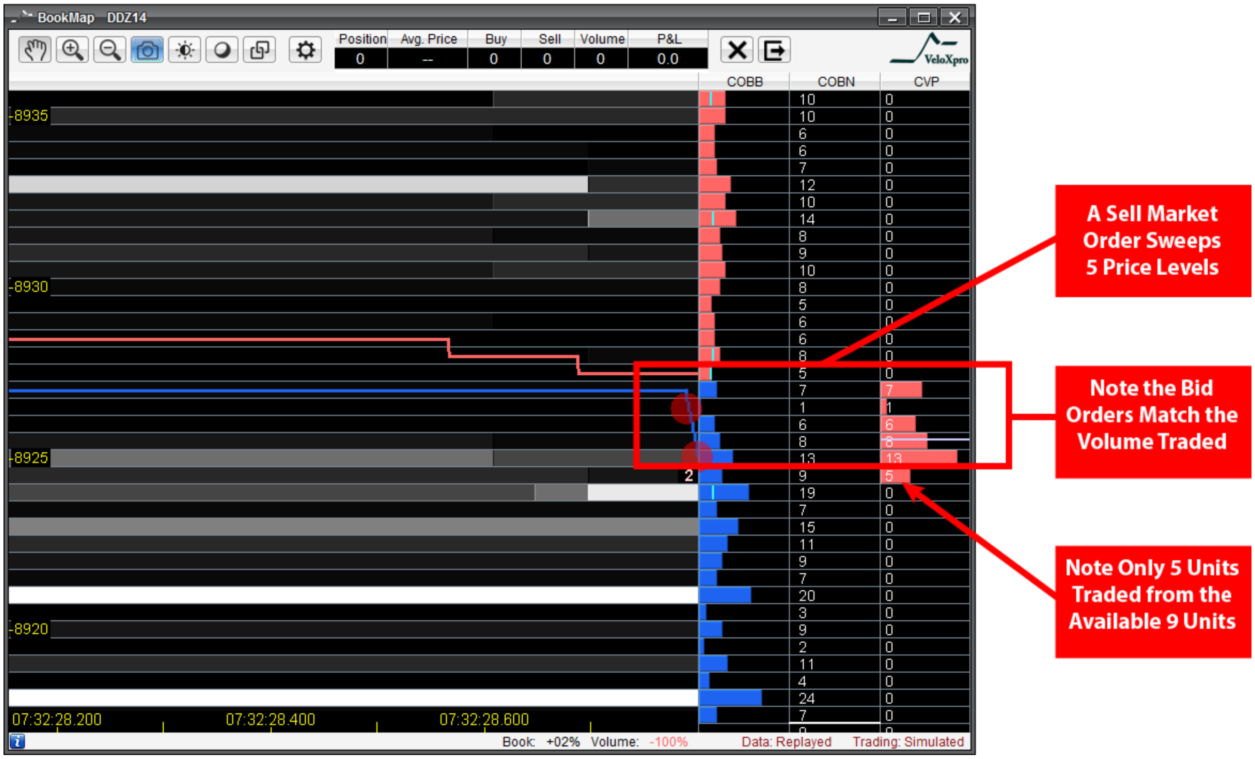

Although most execution venues are based on FIFO, there are some exchanges that possess instruments which utilize pro-rata execution schedules, such as Intercontinental Commodity Exchange (ICE), Chicago Mercantile Exchange (CME), Chicago Board Options Exchange (CBOE), and Philadelphia Stock Exchange (PHLX). With CME's pro-rata schedule, an incoming market buy order is matched with a fixed proportion of each limit order posted at the best ask. Correspondingly, an incoming sell order is matched with equal fractions of all of the limit orders posted at the best bid. Therefore, the larger the limit order at the best bid and the best ask, the larger the fill of that order (Aldridge 2013).

From the exchange point of view, the primary advantage of pro-rata matching is the built-in incentives for traders to place large limit orders which brings liquidity to the exchange. Pro-rata matching motivates traders to post large limit orders without offering special compensation like rebates, and this increases exchange profitability. Additionally, as pro-rata matching eliminates incentives to place and then cancel limit orders with the intent to secure time priority of execution, it reduces message traffic to and from the exchange.

Concisely, the pro-rata incentive works as follows: a trader who wants to execute a limit order is aware that only a fraction of his order will be filled under the pro-rata schedule. As to the size of the filled portion of the order, this will depend on the cumulative size of limit orders placed at the same price by other limit order traders. As the aggregate size of limit orders in a specific price bin goes higher, the percentage of all orders that will be filled in response to an incoming market order of the opposite sign goes lower. And so in order to increase his chances of filling the entire order, the trader is likely to place a limit order with a larger size than his intended order, with the explicit hope that the fraction of the order that will get filled is of the exact size as his intended order.

Figure 2.3 - Shows an example of the pro-rata matching engine. See facing page.


2064.00
2063.75
2063.50
2063.25
Best Offer
2063.00
Best Bid
2062.75
2062.50
2062.25
2062.00
2061.75
100
New Market Sell Order
for 100 hits the Bid

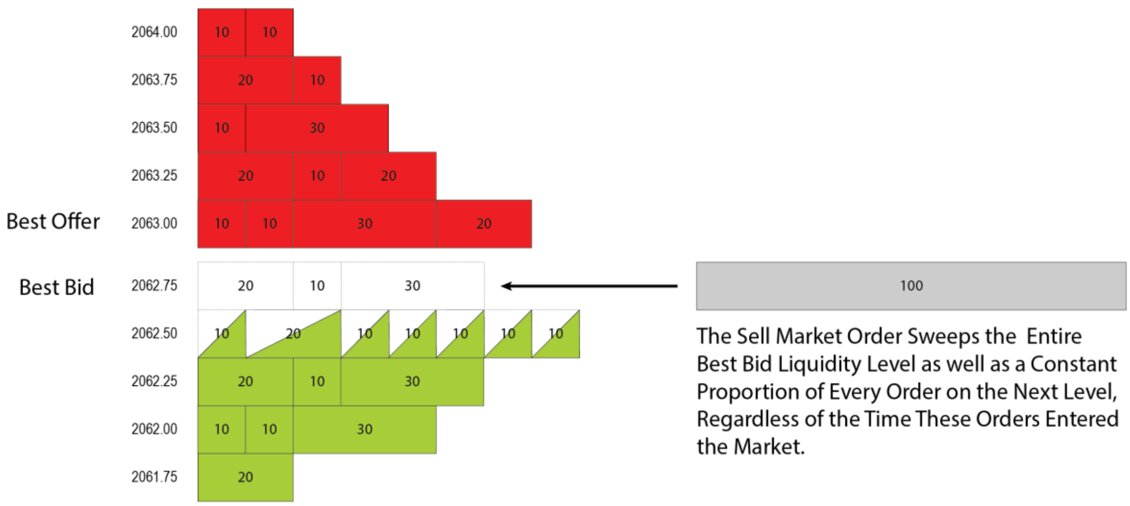
2064.00
2063.75
2063.50
2063.25
Best Offer
2063.00
Best Bid
2062.75
2062.50
2062.25
2062.00
2061.75
100
The Sell Market Order Sweeps the Entire Best Bid Liquidity Level as well as a Constant Proportion of Every Order on the Next Level, Regardless of the Time These Orders Entered the Market.

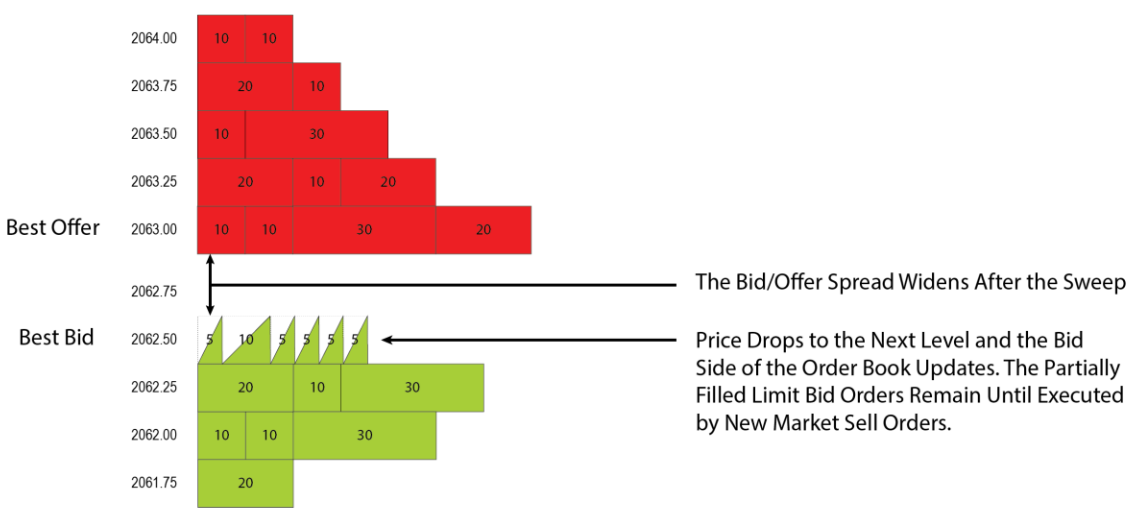
2064.00
2063.75
2063.50
2063.25
Best Offer
2063.00
2062.75
The Bid/Offer Spread Widens After the Sweep
Best Bid
2062.50
2062.25
2062.00
2061.75
Price Drops to the Next Level and the Bid Side of the Order Book Updates. The Partially Filled Limit Bid Orders Remain Until Executed by New Market Sell Orders.

Aggressive versus Passive Order Placement

Orders may be described as either passive or aggressive. Being passive or aggressive indicate how immediately a trader wants to do a trade. Market orders are considered aggressive when the trader wants his order filled right away, at the prevailing market price. Intrinsically, a market order to buy is likely to pay the least offer, while the market order to sell is likely to receive, at the most, the current best bid. On the other hand, a limit order to buy well below the current bid is considered passive. The trader is fine with the low probability of being executed, but when he does execute, he is at least only paying the price he has specified (Narang, 2009).

Aggressive orders do not pertain to malicious orders, and passive orders are not orders that should be taken advantage of. Instead, an order may be passive or aggressive depending on the proximity of its order price to the prevailing market price.

Figure 2.2 - Demonstrates aggressive and passive limit orders.

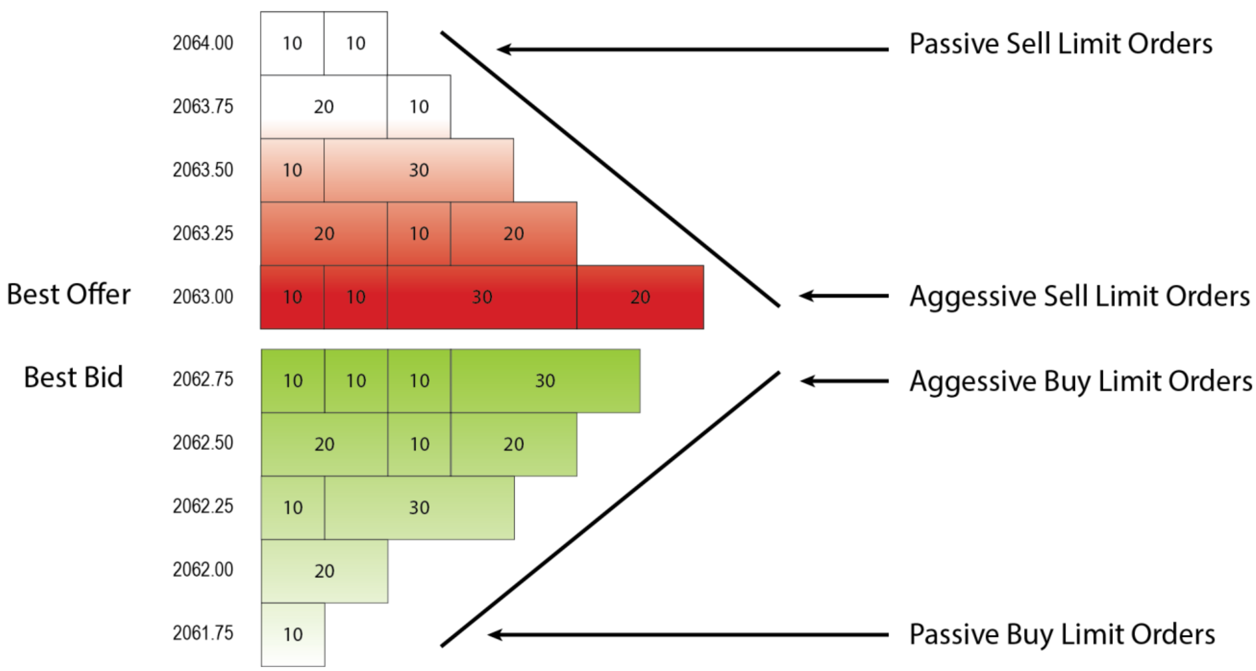

A limit order is considered passive when it is far away from the market price, such as a low-priced limit buy order, or a high-priced limit sell order. When a limit order is closer to the market price, the more aggressive it is deemed to be. The most aggressive order is a market order which "crosses the spread" to be matched with the best-priced limit order on the opposite side of the limit order book. Limit orders crossing the spread are also considered aggressive and are treated like market orders in the execution queue (Aldridge 2013).

Although market orders take advantage of instant and nearly guaranteed fills, market orders cross and pay the spread, incur exchange and broker-dealer transaction fees, and face price uncertainty which, in today's markets, may be the costliest component associated with the market order execution.

Slippage

The market price may "slip" within the time when market order is placed to when the execution is recorded, and this may result in worse execution than the prevailing price at the time the market order was placed. Trade slippage is the cost of missed trades that are never entered or trades that are exited at a worse price because of price limit orders that are not filled (Pardo, 2011). The slippage may be due to several factors:

- Several market orders may have arrived at the exchange and were executed between the time a given market order is placed and executed. Each of these market orders may have depleted the matching liquidity in the order book, adversely affecting the market price. This type of scenario is particularly common at times of news releases, when many traders and their algorithms concurrently process information and place orders in the same direction.

- A market order that is large relative to the available depth of the order book may have swept through the book, executing fractional pieces of the order against limit orders at varying price levels.

- Additional market conditions, such as market disruptions, may also cause significant slippage.

By contrast, the price of a limit order is fixed when the order is placed. A limit order is added to the limit order book, and it "sits" there until the prevailing market price reaches it and a market order is executed against it. Also, limit orders normally avoid "crossing the spread," a cost of paying the market spread incurred by market orders. Highly aggressive limit orders executed as market orders cross the spread, but acquire a good or better execution price than their specified limit price. Limit orders are also subject to positive or negative transaction costs, which differ from one trading venue to another.

Despite their price advantages, limit orders are subject to a critical risk – the risk of non-execution. A limit order is only executed when it is matched with a market order of the opposite direction. Meanwhile, the market price can quickly move away from a limit order, leaving it unexecuted. An unexecuted limit order may present a significant problem when placed to close a position, and miss the opportunity to eliminate market risk of the trade. On the other hand, unexecuted limit orders placed to open a position also incur a cost, that of the opportunity to engage in the trading strategy.

Figure 2.3 – Slope of the Centralized Limit Order Book (CLOB).

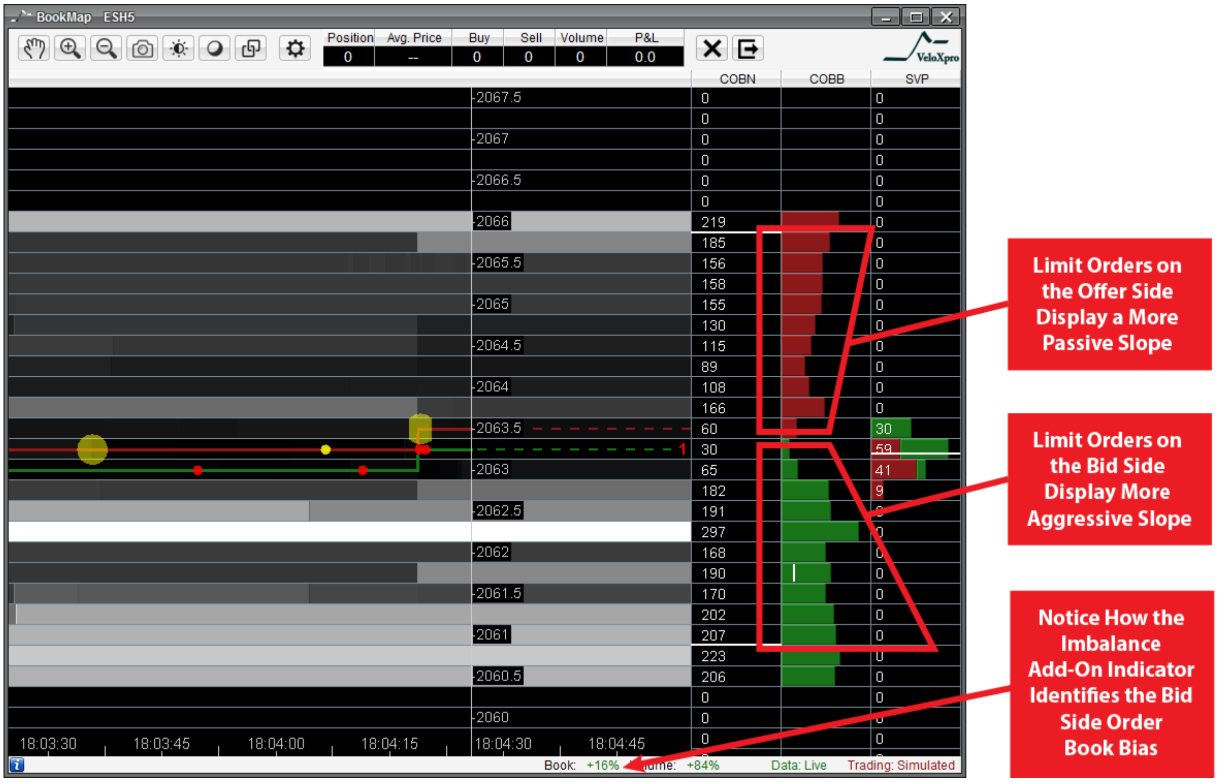

Complex Orders

As competition arises from new entrants in the matching business, trading venues have begun to diversify their order offerings. For instance, competition from dark pools have pushed some exchanges to expand the number of available orders, creating so-called *iceberg orders*. Iceberg orders allow limit-order traders to show only a portion of their order in the limit order book and keep the rest of their liquidity hidden. In FIFO limit order books, iceberg orders are executed on a time priority basis. When matched against a smaller order, the non-executed part of the iceberg is placed back at the end of their limit-order book queue. Unlike the orders in a dark pool, the information as to the size of an iceberg is revealed only after the iceberg is matched in part or in full. The matched size is circulated to other traders as a trade tick. As a rule, iceberg orders cost more than limit and market orders (Aldridge 2013).

Other specialized orders have sprung up to generate additional revenues from higher transaction costs, as well as to cater to the following potential needs of customers:

- *Limited risk*. Most trading venues and broker-dealers now offer a range of orders for controlling market risk, such as hard and trailing stop orders, where the position is liquidated once a price move exceeds the predetermined threshold in the adverse direction.

- *Speed of execution*. Orders in this category offer to deliver the fastest execution possible and include, in addition to the vanilla market order, a market-on-close order which often promises to catch the closing price, a midpoint match order that beats the best bid limit order by endeavoring to negotiate crossing only half of the prevailing spread, and the sweep-to-fill order which instantaneously clears the order book of the size requested in the order. The sweep-to-fill order may be executed faster than the market order, since a large market order is usually executed by sweeping the limit order book progressively over time.

- *Price improvement.* Such orders comprise a block order in options that increases the price by gaining large-volume discounts on transaction costs.

- *Privacy.* Privacy-providing orders, such as iceberg orders and hidden orders, deliver dark liquidity. The iceberg order displays a limited portion of the order in the limit order book, whereas the hidden order is not displayed in the limit order books.

- *Time to market.* Orders in the time-to-market group include fill-or-kill orders for market orders that aim for the fastest liquidity. A fill-or-kill order is canceled if the matching liquidity is not immediately available. Conversely, a good-till-canceled limit order falling in the same order category is kept in the limit order book until it is canceled or another maximum period of time is set by the trading venue.

- *Advanced trading.* These orders include additional quantitative triggers, such as implied volatility in options.

- *Algorithmic trading.* Orders in this category offer execution via order-slicing algorithms.

Figure 2.4 – Diagram of a Hidden Order, or Iceberg Order. See facing page.

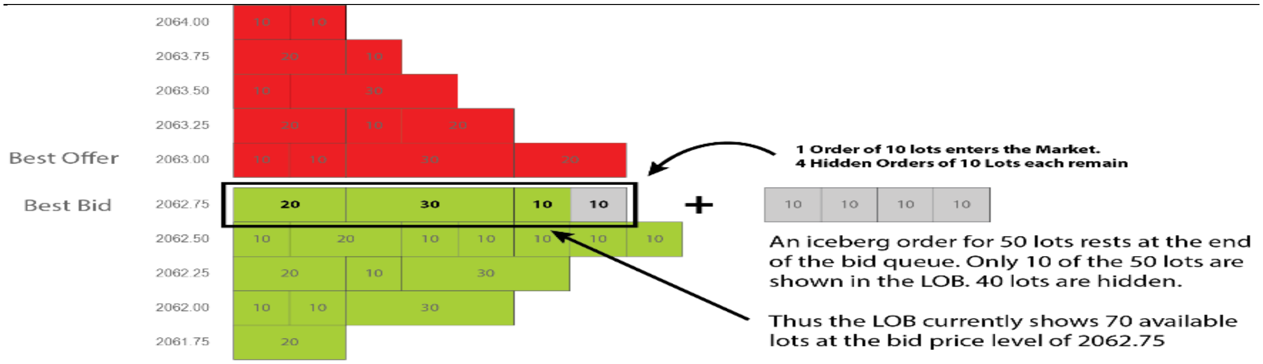
Best Offer
Best Bid
1 Order of 10 lots enters the Market.
4 Hidden Orders of 10 Lots each remain
An iceberg order for 50 lots rests at the end of the bid queue. Only 10 of the 50 lots are shown in the LOB. 40 lots are hidden.
Thus the LOB currently shows 70 available lots at the bid price level of 2062.75

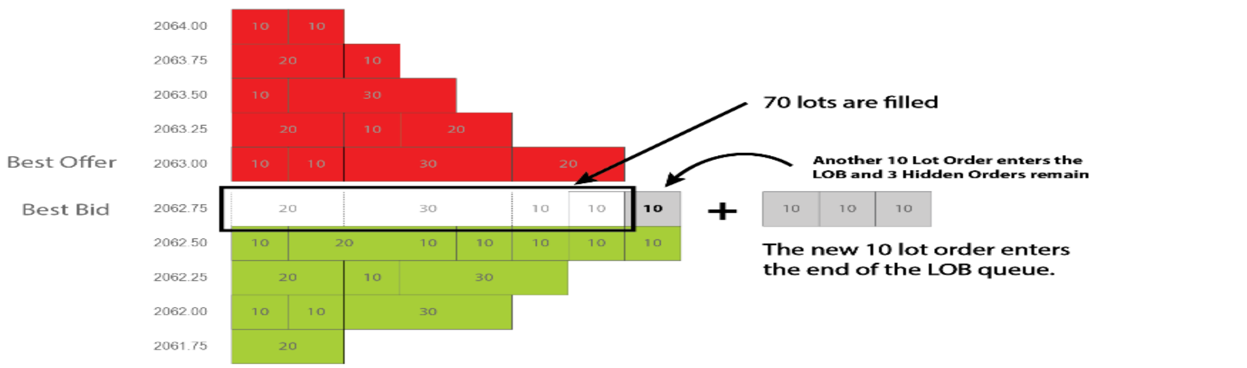
Best Offer
Best Bid
70 lots are filled
Another 10 Lot Order enters the LOB and 3 Hidden Orders remain
The new 10 lot order enters the end of the LOB queue.

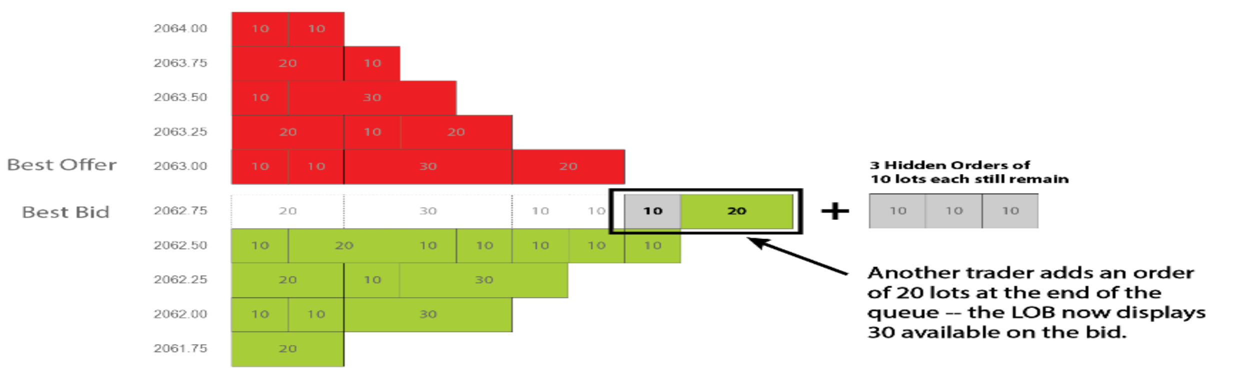
Best Offer
Best Bid
3 Hidden Orders of 10 lots each still remain
Another trader adds an order of 20 lots at the end of the queue -- the LOB now displays 30 available on the bid.

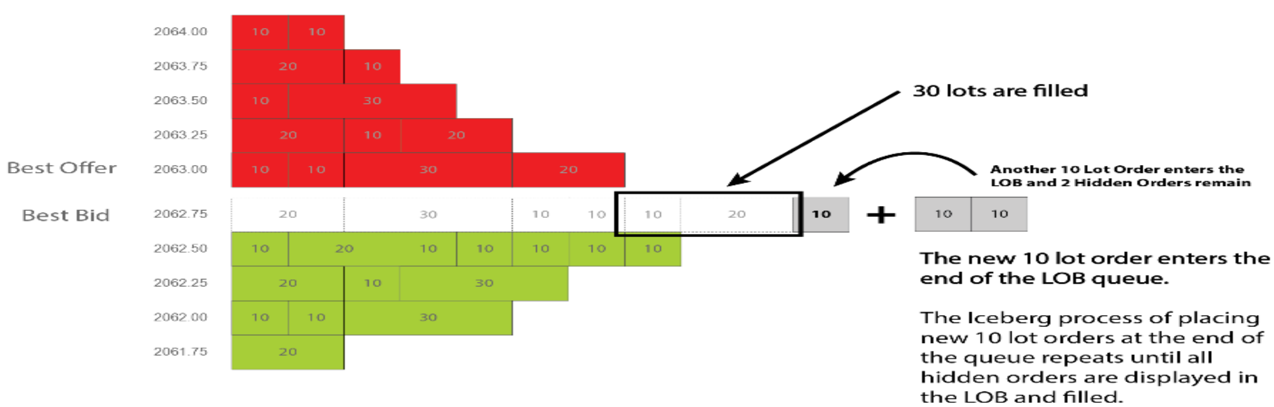
Best Offer
Best Bid
30 lots are filled
Another 10 Lot Order enters the LOB and 2 Hidden Orders remain
The new 10 lot order enters the end of the LOB queue.
The Iceberg process of placing new 10 lot orders at the end of the queue repeats until all hidden orders are displayed in the LOB and filled.

Figure 2.5 – An example of liquidity that trades without a matching limit order at the time of execution.

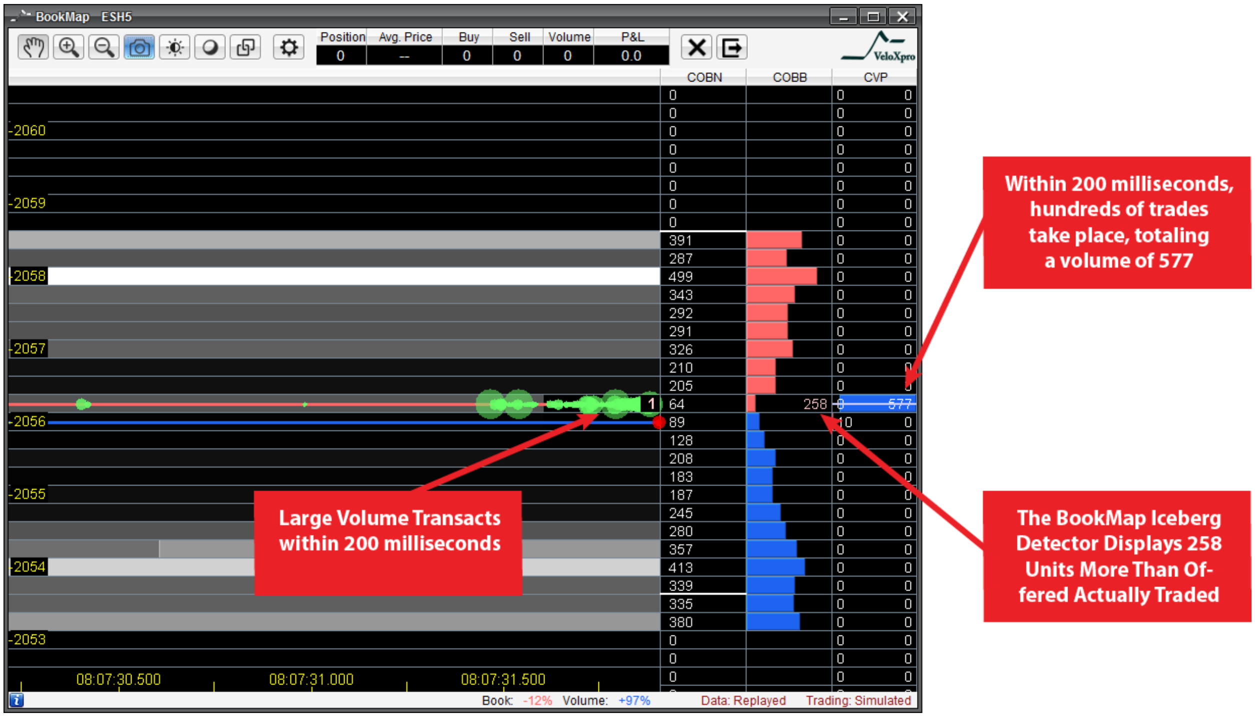

CHAPTER 3: UNDERSTANDING HFT DATA

What is HFT Data?

Trade and quote information is generally distributed in Level I or Level II formats. Level I quotes consist of the best bid price, best ask price, best bid size, best ask size, and last trade price and size, where available. Level II quotes comprise all changes to the order book, such as new limit order arrivals and cancellations at prices away from the market price (Aldridge 2013).

High-frequency data, also known as *tick data*, are records of live market activity. Every time a customer, a dealer, or another entity posts a so-called limit order to buy *s* units of a specific security with ticker X at price *q*, a bid quote is logged at time t_b to buy # units of X.

What happens to quotes from the moment they arrive mostly depends on the venue where the orders are posted. Best bids and asks posted directly on an exchange will be broadcasted to all exchange participants and other parties tracking quote data. In cases when the new best bid is higher than the best ask already in force on the exchange, most exchanges will immediately "match" such quotes, executing a trade at the preexisting best ask, at time t_b. Conversely, if the newly arrived best ask falls below the current best bid, the trade is executed at the preexisting best bid, at time t_a.

Most dark pools match bids and asks by "crossing the spread," but may not broadcast the newly arrived quotes. Likewise, quotes destined for the interdealer networks may or may not be circulated to other market participants, depending on the venue.

Market orders contribute to high-frequency data in the form of "last trade" information. Unlike a limit order which is an order to buy a specified quantity of a security at a certain price, a market order is an order to buy a specified quantity of a security at the best price available at the moment the order is "posted" on the trading venue. As such, market orders are executed immediately at the best available bid or best ask prices, with each market buy order executed at the best ask and each market sell matched with the best bid. The transaction is recorded in the quote data as the "last trade price" and the "last trade size" (Aldridge 2013).

A large market order may require to be matched with one or several best quotes, so several "last trade" data points are generated. For instance, if the newly arrived market buy order is smaller in size than that of the best ask, the best ask quote may remain in force on most trading venues, but the best ask size will be reduced to show that the portion of the best ask quote has been matched with the market order. If the size of the incoming market buy order is greater than the size of the corresponding best ask, the market order consumes the best ask in its entirety, and then proceeds to be matched sequentially with the next available best ask until the size of the market order is fulfilled. The remaining lowest-priced ask quote then becomes the best ask available on the trading venue.

Most limit and market orders are placed in so-called "lot sizes". These are increments of certain number of units, known as a lot. In foreign exchange, a standard trading lot today is US$5 million, a relatively low amount compared to the minimum $25 million entertained by high-profile brokers just a few years back. On equity exchanges, a lot can be as low as one share, however dark pools may still enforce a 100 share minimum requirement for orders. An order for the amount other than an integer increment of a lot size is called "an odd lot."

Small limit and market "odd lot" orders posted via a broker-dealer may be combined or "packaged" by the broker-dealer into larger-size orders to enjoy volume discounts at the orders' execution venue. In the process, the brokers may "sit" on quotes without transmitting them to an executing venue, delaying execution of customers' orders.

How is HFT Data Recorded?

The highest-frequency data are a collection of sequential “ticks”, arrivals of the latest quote, trade, price, order size, and volume information. Tick data usually has the following properties:

- A timestamp
- A financial security identification code
- An indicator of what information it carries:
- Bid price
- Ask price
- Available bid size
- Available ask size
- Last trade price
- Last trade size

Security-specific data, such as implied volatility for options

The market value information, such as the actual numerical value of the price, available volume, or size

A timestamp records the date and time at which the quote is created. It may be the time when the exchange or the broker-dealer released the quote, or the time when the trading system has received the quote. In 2013, the standard “round-trip” travel time in New York of an order quote from the ordering customer to the exchange and back to the customer with the acknowledgment of order receipt is 15 milliseconds or less. Customers have been firing their brokers who were unable to process orders at this now standard speed. Apparently, sophisticated quotation systems include milliseconds and even microseconds as part of their timestamps.

Another part of the quote is an identifier of the financial security. In equities, the identification code can be a ticker. For tickers simultaneously traded on multiple exchanges, a ticker is followed by the exchange symbol. For futures, the identification code consists of the underlying security, futures expiration date, and exchange code.

The last trade price shows the price at which the last trade in the security cleared. Last trade price may differ from the bid and ask. The differences may arise when a customer posts a favorable limit order that is immediately matched by the broker without broadcasting the customer's quote. Last trade size reflects the actual size of the last executed trade.

The best bid is the highest price available for sale of the security in the market. The best ask is the lowest price entered to buy the security at any particular time. Aside from the best bid and best ask, quotation systems may publish “market depth” information – the bid and ask quotes entered posted on the trading venue at prices worse than the best bid and ask, as well as aggregate order sizes corresponding to each bid and ask recorded on the trading venue's “books”.

Market depth information is occasionally referred to as Level II data and may be circulated as a premium subscription service only. On the contrary, the best bid, best ask, last trade price, and size information (“Level I data”) is generally available for a small nominal fee.

How HFT Data is Disseminated

Figure 3.1 - Shows a common use of protocols in trading communication.

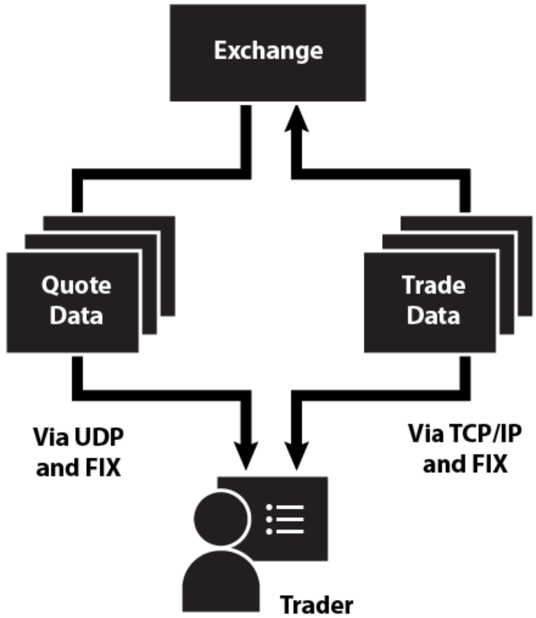

Figure 3.2 - Shows the typical trader data paths.

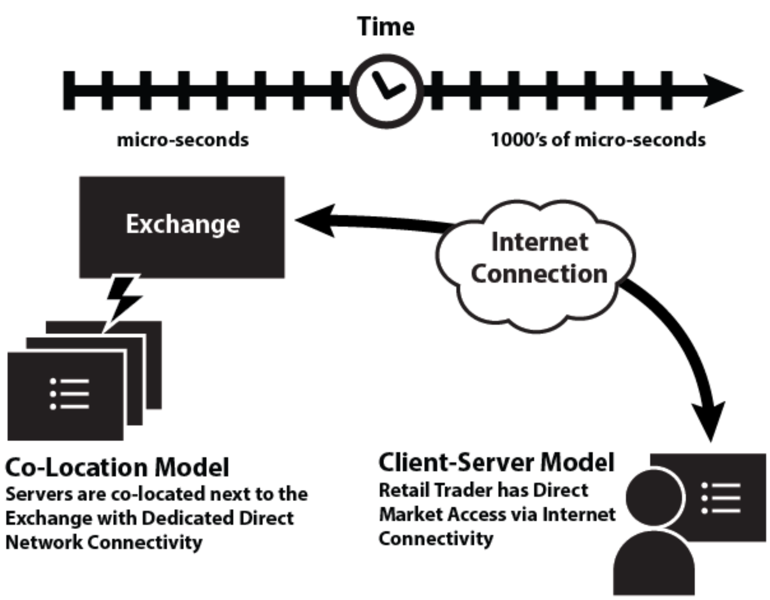

Aggressor Classification

Most Level I or Level II tick data do not contain identifiers that specify whether a given recorded trade was a result of a market buy order or a market sell order, yet some applications call for buy-and-sell trade identifiers as inputs into the models. Some exchanges transmit this data. To address this challenge where the classification is not known or not transmitted from the exchange, or not rebroadcast by a data provider, four methodologies have been proposed to estimate whether a trade was a buy or a sell from Level I data:

- Tick Rule
- Quote Rule
- Lee-Ready Rule
- Bulk Volume Classification (BVC)

The Tick Rule is one of the three most popular methodologies used to ascertain whether a given trade was initiated by a buyer or a seller, when such information is not available in the data set. The other two popular methods are the Quote Rule and the Lee-Ready Rule, after Lee and Ready (1991). The newest method is the Bulk Volume Classification (BVC), due to Easley, Lopez de Prado, and O'Hara (2012).

Based on the tick rule, the classification of a trade is performed by comparing the price of the trade to the price of the preceding trade; no bid or offer quote information is taken into account. Each trade is then classified into one of the four categories:

- *Uptick* - if the trade price is higher than the price of the previous trade.
- *Downtick* - if the trade price is lower than the price of the previous trade.
- *Zero-uptick* - if the price has not moved, but the last recorded move was an uptick.
- *Zero-downtick* - if the price has not moved, but the last recorded move was a downtick.

If the trade's price is different from that of the previous trade, the last trade is classified as an uptick or a downtick, depending on whether the price has moved up or moved down. If the price has not moved, the trade is classified as a zero-uptick or a zero-downtick, depending on the direction of the last non-zero price change.

Figure 3.3 - Shows how BookMap applies the Tick Rule classification method.

Zero Up-tick

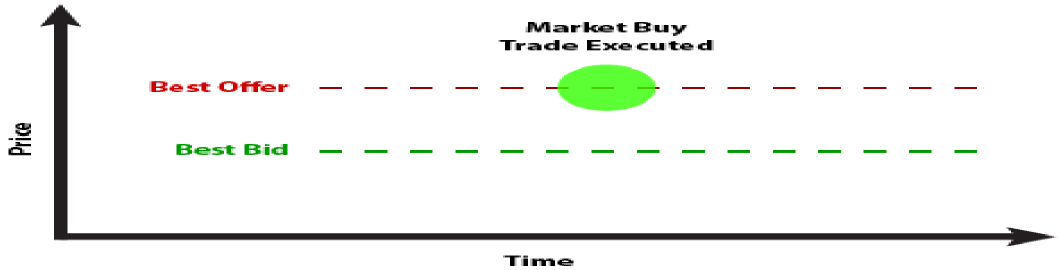

Zero Down-tick

Up-tick

Down-tick

CHAPTER 4: WHAT DOES HFT DO?

What Do High-Frequency Traders Do?

There is no precise definition for HFT, but most market participants agree that HFT strategies fall into the four broad classes below:

1. Arbitrage
2. Directional event-based trading
3. Automated market making
4. Liquidity detection

Figure 4.1 – Shows the four broad HFT "classes"

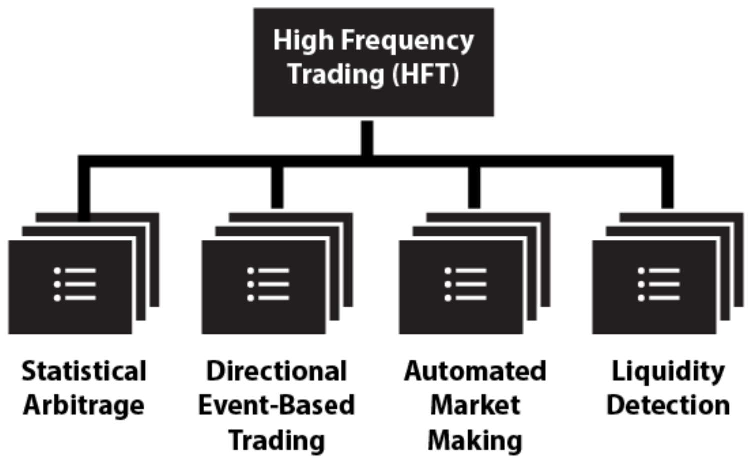

Statistical Arbitrage

Arbitrage strategies trade away price deviations from long-running equilibria or relative asset mispricing. They can include multiple asset classes, as well as multiple exchanges. Many HF arbitrage strategies detect price discrepancies in multiple securities. Several strategies arbitrage prices of the same asset trading on different exchanges, are known as *latency arbitrage strategies*. Most arbitrage strategies are based on assumptions of mean-reversion of asset prices (Aldridge 2013).

Statistical arbitrage models consist of a range of models, including cross-asset models, where financial securities have strong statistical relationships. These models are deep-rooted in economic theories, ruling out spurious statistical relationships developed using plain data mining and also known as the Spaghetti Principle of Modeling which states that if one throws a plate of spaghetti filled with data against the wall of statistics, something may stick. However, what sticks may not have any sound reason for sticking and is likely to fall apart in production.

Case in point, bonds and interest rate futures have been shown to possess substantial dependencies and, as a result, their values tend to move in tandem. When prices of bonds or interest rate futures digress from their long-running price equilibrium for no palpable reason, statistical arbitrage trading may be feasible by purchasing the instrument with a lower-than-expected price relative to the other instrument(s), and selling the instrument with a higher-than-expected price relative to the other instrument(s).

Figure 4.2 – Shows a simplified framework for a Statistical Arbitrage type strategy.

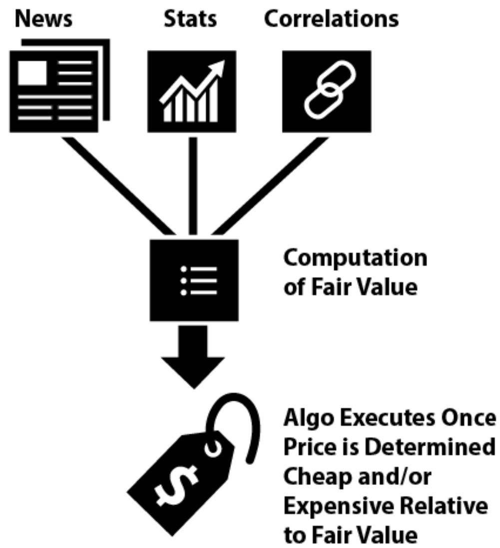

Directional

Directional strategies identify short-term trend or momentum. This class of high-frequency strategies includes event-driven strategies, strategies based on predictable short-term price movements, as well as the controversial ignition strategies. Event arbitrage models show the methodology as well as performance of trading on predictable and recurrent effects of news. For some, access to the news data is paramount to the strategy. See: http://www.cnbc.com/id/100809395 and http://www.wsj.com/articles/fast-traders-are-getting-data-from-sec-seconds-early-1414539997for more background

Figure 4.2 – Shows an example of news event directional Event-Based Trading.

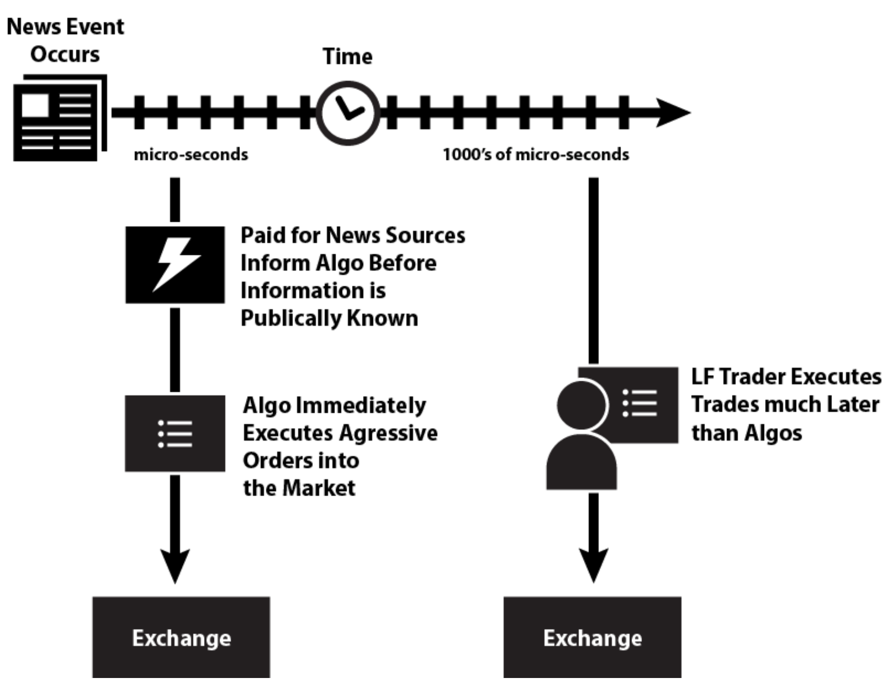

Automated Market-Making

Automated market-making strategies include perhaps the most traditional trading strategies, such as automated market-making, a cost-effective and accurate alternative to human broker-dealers. The category of automated market-making and liquidity provision comprises both inventory-driven and information-driven approaches. Inventory-driven methods are inclined to focus on joint minimization of the inventory risk and market risk, guaranteeing that the positions are within a trader's risk tolerance limits given market conditions, and evaded where appropriate (Aldridge 2013).

Information-driven market-making models are built with the purpose of curtailing the risk of adverse selection, the risk of taking an opposite position to a better-informed party. To reduce the number of such losing positions, high-frequency traders can organize a wide range of models that will assist in forecasting short-term directionality of markets, tracking the number of well-informed market players in the market waters, and predicting imminent lumps and shortages of liquidity. These techniques allow traders to select the quantities and levels of aggressiveness of their orders based on expectations of surplus or shortage of liquidity.

Figure 4.3 – Shows how BookMap identifies algorithmic activity within the HFT environment. See facing page.

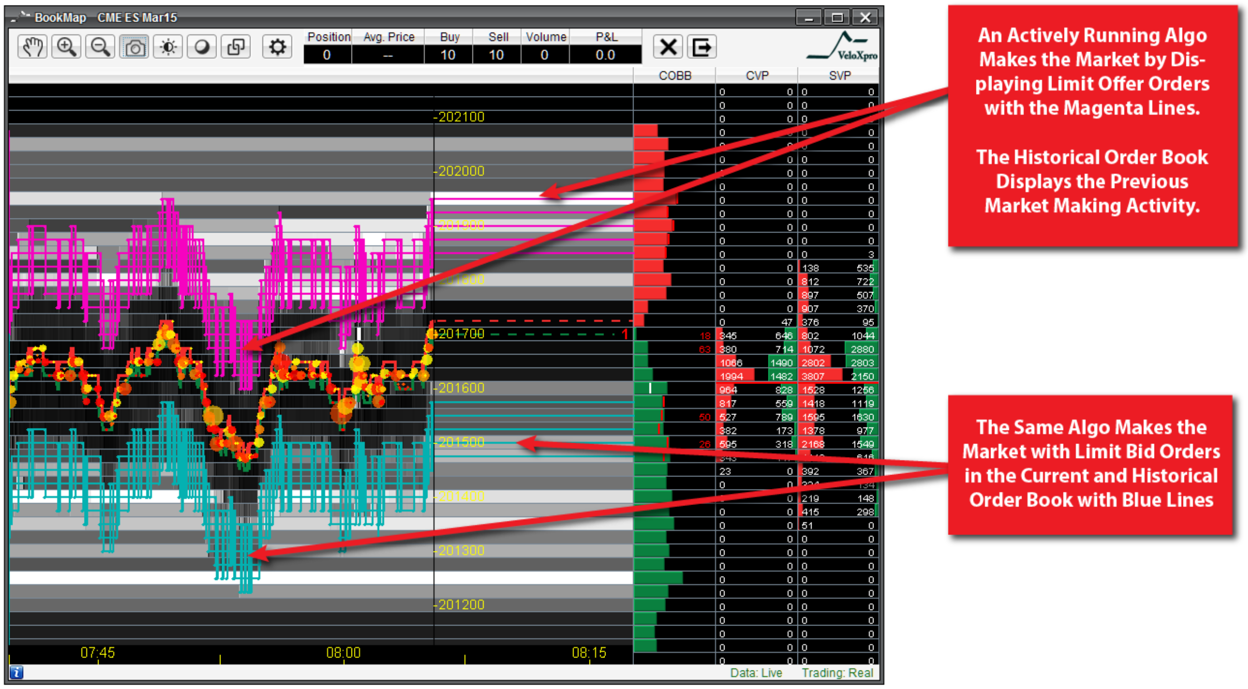
BookMap - CME ES Mar15
An Actively Running Algo Makes the Market by Dis-playing Limit Offer Orders with the Magenta Lines.
The Historical Order Book Displays the Previous Market Making Activity.
The Same Algo Makes the Market with Limit Bid Orders in the Current and Historical Order Book with Blue Lines

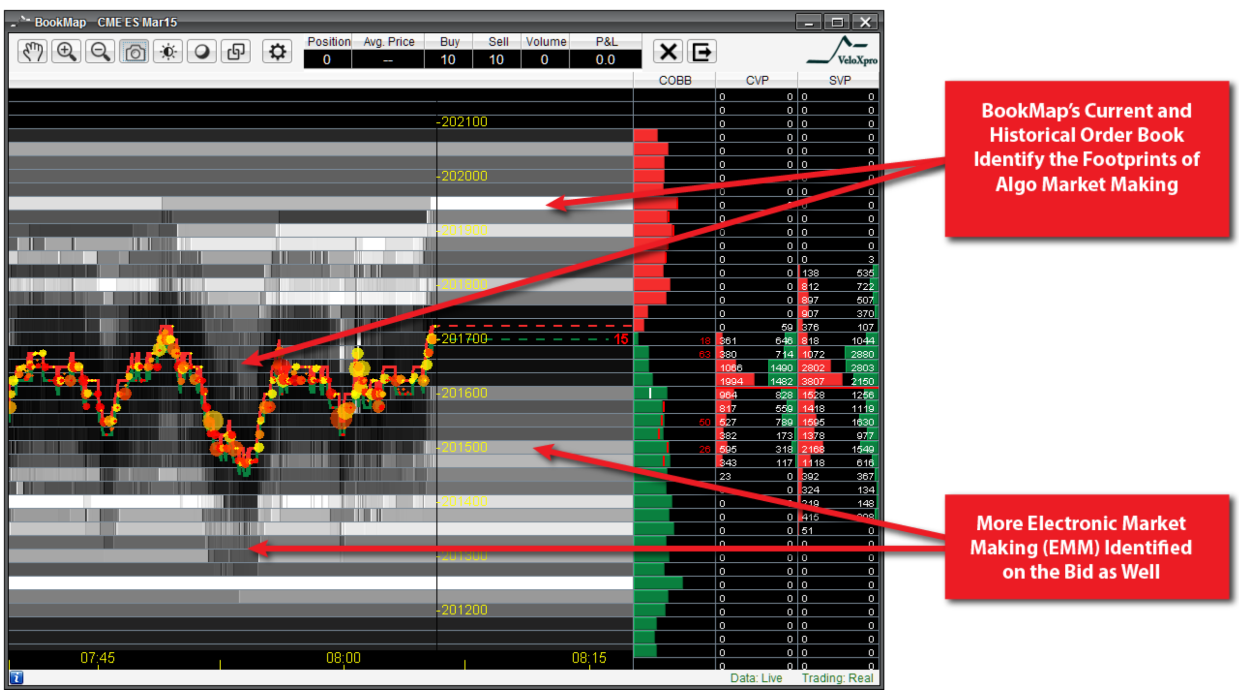
BookMap - CME ES Mar15
BookMap's Current and Historical Order Book Identify the Footprints of Algo Market Making
More Electronic Market Making (EMM) Identified on the Bid as Well

Liquidity Detection

Liquidity detection strategies, such as pinging (also known as sniffing and sniping), quote stuffing, and spoofing, may perhaps be the least palatable to the low-frequency investors. “Pinging” has been detected to exist on selected venues, such as dark pools. In ping orders, small orders are entered in order to determine the level of hidden orders. This strategy is particularly used to evaluate what is resting on a dark platform. Quote stuffing is when large numbers of orders and/or cancellations/updates to orders are entered to create uncertainty for other participants and slow down their process. It is also intended to conceal their own strategy. Layering or spoofing involves the submission of multiple orders, usually away from the touch of one side of the order book, to execute a trade on the other side of the order book so that once the trade has taken place, the manipulative orders would be removed. Momentum ignition is the entry of orders or a series of orders with the intent to start or intensify a trend, as well as to encourage other participants to fast-track or extend a trend so that there is an opportunity to unwind or open a position at a favorable price. The nature of other strategies like “ignition strategies” have been mostly hypothetical, and no reliable evidence of strategy existence has been produced to date (Gregoriou, 2015). Contrary to some opinions, a clear example of what appears to be an “ignition” algorithm in action is shown here: https://goo.gl/x0rTwC and here: https://goo.gl/P1FNlP

A great reference for further research is “High frequency trading, by Bruno Biais (Toulouse School of Economics)” 2011

Figure 4.4 – Displays how some algos entice retail traders with dangling orders.

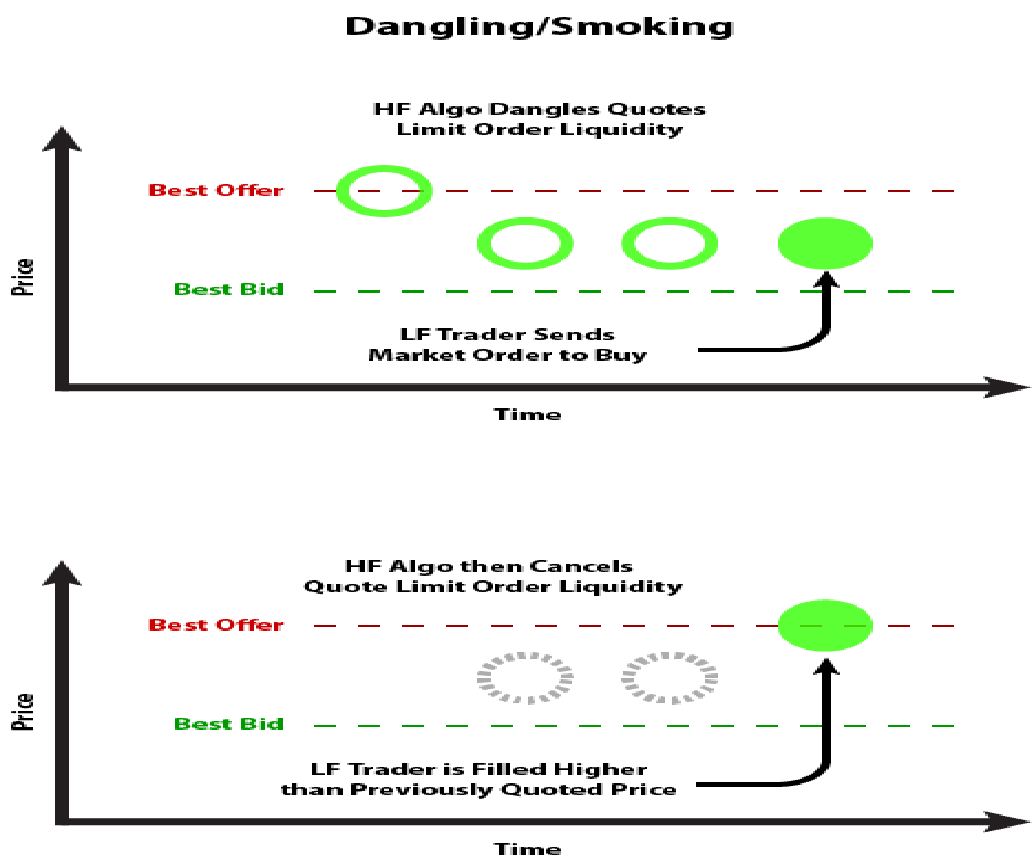

Figure 4.5 – Shows a definition of spoofing

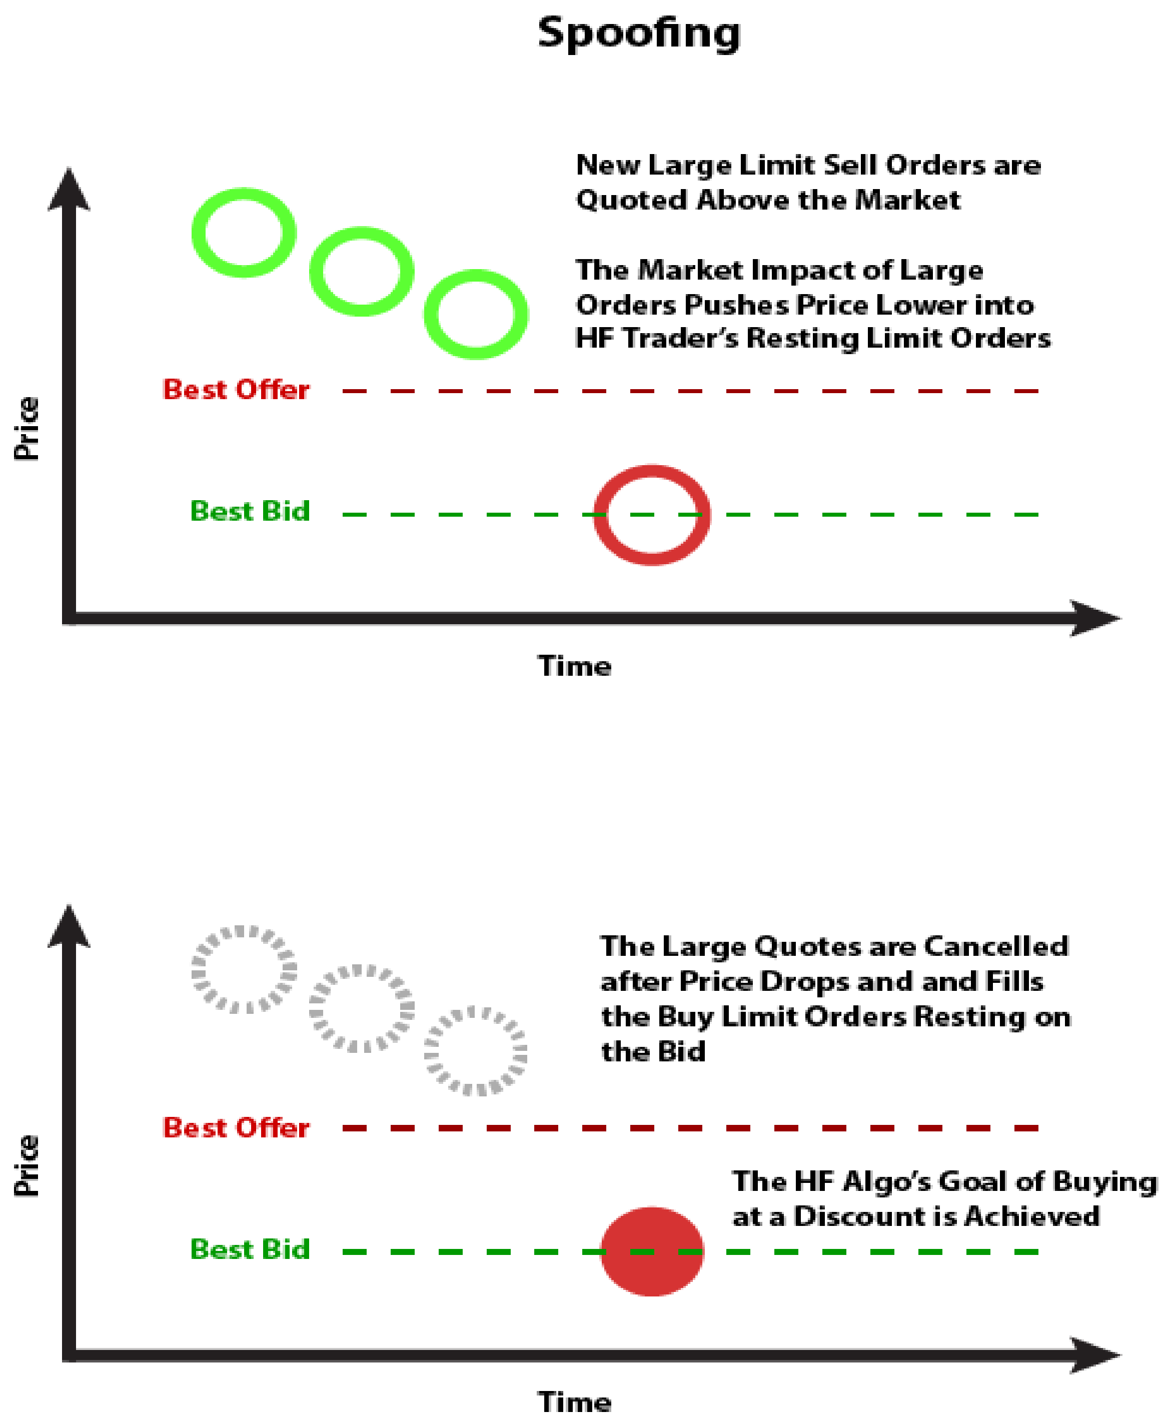

Figure 4.6 – Shows how HFT algos use stuffing schemes to flood the exchange with quotes and orders.

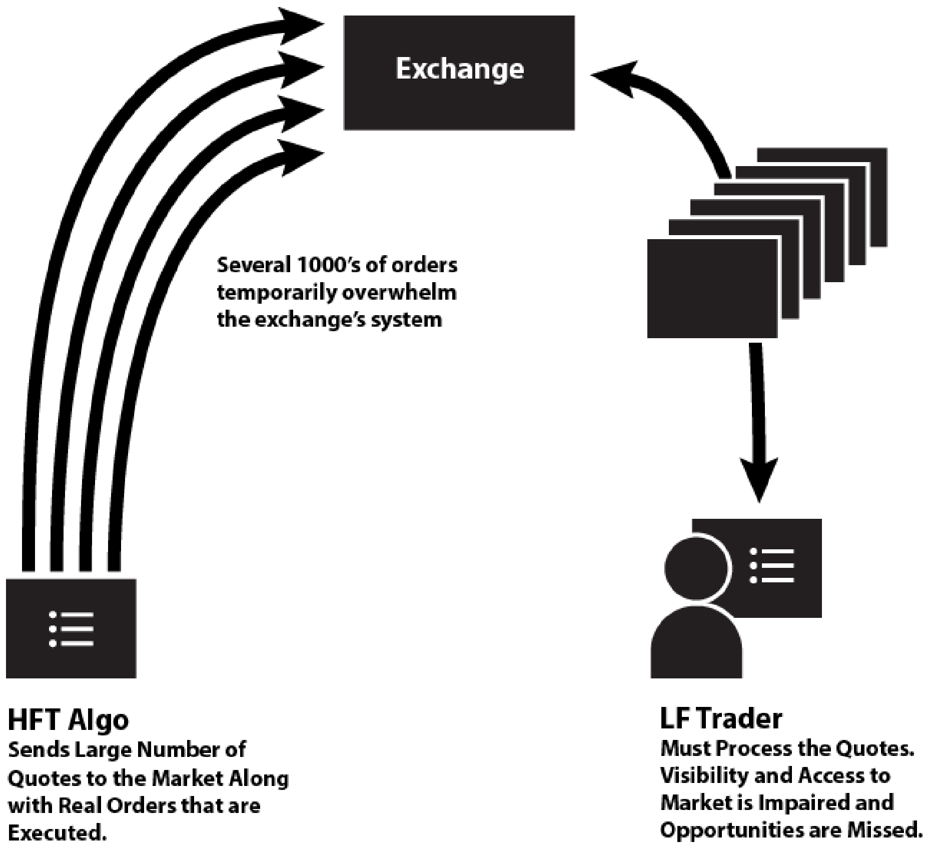

Another great resource for this information is found here: https://goo.gl/WjAAp6

BIBLIOGRAPHY

Aldridge, Irene. *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. 2nd ed. New York: John Wiley & Sons, 2013.

Hendershott, Terrence; Jones, Charles; Menkveldf, Albert.*Does Algorithmic Trading Improve Liquidity?The Journal of Finance.Volume 66 Issue 1*. 2011. [Online] Available from: http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.2010.01624.x/full [Accessed 23 April 2015].

Pardo, Robert. *The Evaluation and Optimization of Trading Strategies.*New York: John Wiley & Sons, 2011.

Gregoriou, Greg. *Handbook of High Frequency Trading.*Academic Press, 2015.

Narang, Rishi. *Inside the Black Box: The Simple Truth About Quantitative Trading*.
New York: John Wiley & Sons, 2009.

Biais, Bruno (Toulouse School of Economics)*High frequency trading: Presentation prepared for the European Institute of Financial Regulation.*Paris, Sept 2011 [Online] Available from: http://www.eifr.eu/files/file8379010.pdf [Accessed 23 April 2015]