







Unravelling the JPMorgan spoofing case using particle physics visualization methods

Philippe Debie^{1,2}  | Cornelis Gardebroek³  |
 Stephan Hageboeck⁴  | Paul van Leeuwen⁵ |
 Lorenzo Moneta⁶ | Axel Naumann⁶  |
 Joost M. E. Pennings^{1,7,8,9} | Andres A. Trujillo-Barrera¹⁰  |
 Marjolein E. Verhulst^{1,11} 

¹Marketing and Consumer Behaviour Group, Wageningen University, Wageningen, The Netherlands

²Innovation- and Risk Management and Information Governance, Wageningen Economic Research, Den Haag, The Netherlands

³Agricultural Economics and Rural Policy Group, Wageningen University, Wageningen, The Netherlands

⁴Information Technology Department, European Organization for Nuclear Research (CERN), Geneva, Switzerland

⁵Data Science, Information Management & Projectmanagement Organisation, Wageningen Economic Research, Den Haag, The Netherlands

⁶Experimental Physics Department, European Organization for Nuclear Research (CERN), Geneva, Switzerland

⁷Department of Marketing and Supply Chain Management, Maastricht University, Maastricht, The Netherlands

⁸Department of Finance, Maastricht University, Maastricht, The Netherlands

⁹Office for Futures and Options Research, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA

¹⁰Department of Agricultural Economics and Rural Sociology, University of Idaho, Moscow, Idaho, USA

¹¹Consumer and Chain, Wageningen Economic Research, Den Haag, The Netherlands

Correspondence

Marjolein E. Verhulst, Wageningen University & Research, Hollandseweg 1, 6706 KN, Wageningen, The Netherlands.
 Email: marjolein.verhulst@wur.nl

Abstract

On 29 September 2020, JPMorgan was ordered to pay a settlement of \$920.2 million for spoofing the metals and Treasury futures markets from 2008 to 2016. We examine these cases using a visualization method developed in particle physics (CERN) and the messages that

We would like to thank the Editor (John A. Doukas) and two anonymous reviewers for their constructive comments and suggestions that significantly improved the quality of this article.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. *European Financial Management* published by John Wiley & Sons Ltd.

the exchange receives about market activity rather than time-based snapshots. This approach allows to examine multiple indicators related to market manipulation and complement existing research methods, thereby enhancing the identification and understanding of, as well as the motivation for, market manipulation. In the JPMorgan cases, we offer an alternative motivation for spoofing than moving the price.

KEYWORDS

high-frequency trading, limit order book, particle physics, spoofing, visualization

JEL CLASSIFICATION

G10, G18, G23, G28, K22, K23

1 | INTRODUCTION

‘A LITTLE RAZZLE DAZZLE TO JUKE THE ALGOS...’ wrote a JPMorgan Treasury trader in a chat message in November 2012, after successfully tricking high-frequency traders and moving the market (Schoenberg & Robinson, 2020). Fast forward to the year 2020, and JPMorgan (JPM) had to pay a record-breaking settlement of \$920.2 million for manipulating the precious metals and Treasury markets (Commodity Futures Trading Commission [CFTC], 2020; Michaels, 2020). Specifically, JPM¹ admitted to spoofing the gold, silver, platinum, palladium, Treasury note and Treasury bond futures markets² between 2008 and 2016.

Spoofing has been illegal under the Dodd-Frank Act since 2010 and is defined as: ‘*bidding or offering with the intent to cancel the bid or offer before execution*’ (United States, 2010, p. 1739). Spoofers manipulate the displayed order volume³ (hereafter referred to as ‘order volume’) in the limit order book (LOB) to persuade market participants to trade in the spoofer’s desired direction (Dalko & Wang, 2018). The LOB shows the order volume at various price levels. However, it presents incomplete information to market participants. For example, market participants do not know what type of order is submitted, the actual volume of an iceberg order and whether a reduction in volume is due to a cancellation or an order execution (Dalko & Wang, 2018). Spoofers can take advantage of this market microstructure by introducing conditions that can influence the decisions of other traders (Mendonça & de Genaro, 2020).

One of the basic types of spoofing involves the spoofer wanting to buy at a lower price than the current price (Dalko & Wang, 2018): a relatively small genuine order (i.e., an order intended to be

¹JPMorgan Chase & Company and its subsidiaries, JPMorgan Chase Bank and J.P. Morgan Securities LLC.

²These futures contracts were/are traded on the Commodity Exchange, Inc. (COMEX), the New York Mercantile Exchange (NYMEX) and the Chicago Board of Trade (CBOT).

³Contrary to hidden order volume, which can be the case with iceberg orders, an iceberg order is an order whereby only a fraction of the total order is displayed in the LOB and the rest is not visible to other market participants (Buti & Rindi, 2013).

executed) is placed on the bid side and a relatively large spoof order (i.e., an order not intended to be executed) is placed on the opposite side—the ask side—of the LOB. Market participants then act on the newly created imbalance in the LOB and move the market in the direction of the genuine order's price, often by way of herd behaviour (Dalko & Wang, 2018). Shortly after placing the spoof order, or once the genuine order has been executed, the large spoof order is cancelled and the imbalance created is gone. The result is that the spoofer was able to buy at a lower price (CFTC, 2020; Dalko & Wang, 2018). Other types of spoofing include, but are not limited to, layered spoofing, layered spoofing with collapsing and spoofing with vacuuming and flipping (Neurensic, 2016). Spoofing can be hard to identify as it may, for example, take place within a single market, between correlated markets (e.g., soybean futures and soybean oil futures), between different calendar contracts (e.g., the March and September contracts of E-mini S&P 500 futures), between derivatives (e.g., gold futures and gold options), between exchanges and by one party or by multiple parties. Moreover, spoofing concerns the trading intention to cancel before execution, and 'intention' is difficult to capture in market data.

Spoofing is harmful to markets and their participants for numerous reasons. Spoofers intentionally distort the available information that traders use to make decisions. This makes nonspoofing market participants vulnerable as they are misguided by false buy and/or sell liquidity figures (Dalko & Wang, 2018). This negatively impacts the price formation process and hence distorts the price (Dalko & Wang, 2020; Mendonça & de Genaro, 2020). It also creates additional volatility in price, trading volume and order volume, which negatively impacts the stability of the market (Dalko & Wang, 2020). Moreover, its effects can spill over into interconnected markets, making them inefficient too (Mendonça & de Genaro, 2020).

Over the course of 8 years, JPM placed hundreds of thousands of spoof orders resulting in \$172,034,790 in gains. Conversely, however, these orders harmed the market and its participants, causing \$311,737,008 in market losses (CFTC, 2020). As this only represents identified spoofing by one firm, the real damage caused by spoofing across all markets is likely to be much greater, making this a serious problem for all stakeholders. The current supervisory systems are not adequate and effective enough to detect such illegal trading behaviour, given that (1) JPM's supervision system failed to detect manipulative practices, such as spoofing, until 2014 (CFTC, 2020) and (2) it took the CFTC 3–11 years after the spoofing occurred to file charges against JPM and many of the spoofing instances were probably discovered thanks to secured documents and computer communication.

Using a visualization methodology developed in particle physics by the European Organization for Nuclear Research (CERN) (Antcheva et al., 2009; CERN, 2018b; Verhulst et al., 2021), we describe the LOB in a novel way, providing new insights into the JPM spoofing case. Specifically, we visualize all spoofing examples as documented in the CFTC report (CFTC, 2020). It contributes to the literature as follows. First, we offer guidance on how to characterize spoofing by way of variables and how to effectively visualize these variables. Second, we offer an alternative motive for spoofing, namely, attracting liquidity rather than changing the price. To the best of our knowledge, this has not been reported before. Third, we provide insight into how spoofing is conducted and how (in)visible it is to other market participants. Fourth, while previous LOB visualizations were solely time-based (i.e., using snapshot intervals of, e.g., 5 seconds), this study complements these visualizations with the original messages about traders' market activity as sent to the exchange. Moreover, it illustrates how high-frequency LOB data can be effectively visualized. This novel way of visualizing high-frequency data can contribute to new insights in future research and inspire further analyses among stakeholders. Companies such as JPM, for example, can use the methodology to enhance and refine their surveillance programs and internal control systems, and regulators, such as the CFTC, can use it to enhance their understanding of manipulative trading practices.

2 | LITERATURE REVIEW

Empirical literature on spoofing is scarce, particularly due to constraints in obtaining (LOB) data that matches the purpose of the research (Lee et al., 2013; Linton & Mahmoodzadeh, 2018; Putniņš, 2012). Several studies have tried to detect spoofing in markets by using order data (Lee et al., 2013; Zhai et al., 2017). This data differs from LOB data, in that order data comprises the submitted, cancelled and modified orders of individual traders, whereas LOB data constitutes all these orders and shows the LOB visible to all market participants. For example, LOB data reveals the best bid and ask prices and total volumes belonging to specific price levels in the LOB (Mendonça & de Genaro, 2020). Although order data contains more information on individual orders (provided that it is not aggregated), studies attempting to detect spoofing with order data have omitted to reconstruct the LOB, let alone visualize it.

LOB data is nevertheless needed to understand the current state of the market, which influences trading and spoofing decisions. It helps to identify higher-level patterns or parameters related to spoofing, for example, imbalances between the bid and ask volumes (Cartea et al., 2020). To the best of our knowledge, there are only a handful of researchers who studied spoofing using LOB data. Mendonça and de Genaro (2020) generated 1-min LOB snapshots from order data and used both datasets to detect spoofing on the Brazilian Stock Exchange. Leangarun et al. (2016) tried to detect, among others, spoofing in three NASDAQ stock markets by training neural networks and using 1-min LOB intervals.

However, these papers, as well as other LOB-related papers (e.g., Biais et al., 2010; Menkveld & Yueshen, 2019), lack visualizations of the LOB. Visualizations help academics, industry participants and regulators to better understand the market; they allow them, among other things, to identify and understand anomalies such as spoofing (Verhulst et al., 2021). LOB visualization literature is thus scarce (Aidov & Daigler, 2015; Paddrik et al., 2016). In addition, visualizations that do exist are often time-based and thus have limitations: because orders arrive irregularly, order data and LOB data are irregularly spaced over time. To achieve regular time intervals, time-based visualizations and time-series analyses use snapshots of the LOB. As a result, information is lost as the information is being aggregated. In addition, the literature provides no uniform method to achieve optimal snapshot size (Verhulst et al., 2021). Snapshot sizes that have been used in LOB analyses so far are: 5 minutes (Chordia et al., 2019; Kahraman & Tookes, 2017), 1 minute (Hautsch & Horvath, 2019; Mendonça & de Genaro, 2020; Yao & Ye, 2018), 10 seconds (Cont et al., 2014), 5 seconds (Brogaard & Garriott, 2019), 3 seconds (Ito & Yamada, 2018) and 1 second (Battalio et al., 2016; Brogaard et al., 2019; Colliard & Hoffmann, 2017; Dugast, 2018). This lack of uniformity can be explained by the ever-increasing velocity (size) of data. At the start of the 21st century, one day of message data was comparable in size to 30 years of daily data (Dacorogna et al., 2001). Ten years later, data velocity had increased tenfold (Fabozzi et al., 2011). With today's high-frequency orders, 1-s intervals can contain thousands of orders and action/reaction cycles of algorithms, hence increasing the need for a high resolution. Past research has identified the benefits of high-frequency trading (HFT) for market participants. Brogaard (2010) shows that HFT adds substantially to the price discovery process and Brogaard et al. (2014) find that HFT facilitates price efficiency by trading in the direction of permanent price changes and in the opposite direction of transitory pricing errors. Hasbrouck (2018) examines high-frequency quoting and finds, among others, a positive relation between competition and quote volatility. He indicates that his analysis is directed at a broad classification of quote volatility and does

not rule out occurrences of quote stuffing or spoofing (Hasbrouck, 2018, p. 636). Here, we exclusively focus on spoofing as an example of HFT, and as such, this paper may contribute to literature that examines the role and impact of HFT on financial markets. Moreover, it complements past literature on LOBs and existing visualizations by applying visualization methodologies from particle physics to message-based LOB data.

3 | DATA AND METHODOLOGY

Data consists of the Chicago Mercantile Exchange (CME) Group's proprietary market-depth data set for all spoofing examples reported by the CFTC (CFTC, 2020). The files are in the CME Market Depth 3.0 format, which provides messages about market activity.⁴ These messages can be used to recreate the LOB with millisecond precision. The open-source ROOT software framework, developed by CERN, among others, to analyse the massive data generated in the Large Hadron Collider, is used to reconstruct and visualize the LOB (Brun & Rademakers, 1997; CERN, 2018b; Verhulst et al., 2021). ROOT is used in particle physics to save, access and mine data, among other applications, as well as to generate visualizations (CERN, 2018a). Large amounts of data can be stored and processed efficiently in a distributed setup (Tejedor & Kothuri, 2018).

The CFTC (2020) reported nine specific examples of spoofing and manipulation by JPM, including the associated markets, dates, timestamps (Central Time), volume orders and prices. We discuss the nine examples according to their spoofing strategies: (1) 'traditional' spoofing, that is, there is a displayed genuine order and a single spoof order; (2) spoofing with iceberg orders, that is, the genuine order is an iceberg order with displayed and hidden volumes and a single spoof order; (3) layered spoofing, that is, there is a displayed genuine order and multiple spoof orders at various price levels and (4) layered spoofing with iceberg orders, that is, the genuine order is an iceberg order with displayed and hidden volumes and there are multiple spoof orders at various price levels. Section 4 discusses only one example per spoofing category, and meaningful differences will be noted. Figures and tables for all spoofing examples not discussed in this paper are available in the Online Supporting Information Appendix.

Time windows in which the spoofing examples took place are visualized using ROOT's graphing facilities (Brun & Rademakers, 1997; CERN, 2018b). For readability, only the top 10 bid and ask levels are visualized from the consolidated limit order book.⁵ First, we will show the LOB for a single spoofing example using two snapshot sizes employed in previous literature: a 5-s snapshot (Brogaard & Garriott, 2019) and a 1-s snapshot (Battalio et al., 2016; Brogaard et al., 2019; Colliard & Hoffmann, 2017; Dugast, 2018). Subsequently, we visualize the LOB message by message, rather than time-based visualizations which are customary in the existing literature. We use messages as they have (almost) the same data granularity as trading algorithms, and we demonstrate that these visualizations show what is actually happening in the market. Second, we highlight one spoofing example for each spoofing category by enriching the visualizations with variables that may further characterize spoofing behaviour. We provide

⁴The data contains messages on LOB level changes (i.e., a new price level is inserted or deleted, or the volume is changed at a level). If two traders add volume at the same price level in the same millisecond, for example, there will be a single message about the aggregated volume addition, rather than two separate messages (i.e., one for each trader).

⁵Several LOBs from the spoofing examples contain more than 10 levels because of the implied LOB, but these are not visualized as they are generally further away from the top 10 bid and ask levels. Visualizing all levels would make the visualizations unreadable in a paper version.

a unique visualization of the LOB in the relevant time window, showing: (1) the prices and volumes of all LOB levels; (2) midpoint prices; (3) the number of messages received by the exchange; (4) cumulative trade volume and individual trades including their respective prices; (5) volumes of the first bid and ask levels; (6) cancelled volume on the first bid and ask levels and (7) bid and ask side liquidity.

Liquidity is measured by the Adverse Price Movement (APM) of the Exchange (or Xetra) Liquidity Measure (Gomber & Schweickert, 2002; Gomber et al., 2015; Sensoy, 2019). APM bid (APM ask) represents the execution costs in basis points (bps) of a trader who immediately wants to sell (buy) a dollar value and takes liquidity from the bid (ask) side by submitting market orders. A lower APM indicates that the cost of trading is low and, therefore, liquidity is high (Gomber & Schweickert, 2002). For each message, the total LOB dollar value is calculated by multiplying the LOB prices with their respective volumes. The mean dollar value is calculated for the respective month in which the spoofing example took place and is used for the APM calculation. To test whether significant changes in liquidity occur before, during and after spoofing, the data is split into three parts for each spoofing example: 'before' represents the time up until the spoof order was added; 'during' the period from when the spoof order was added until it was cancelled and 'after' the time following the cancellation of the spoof order. Five different time windows are used: (1) the same time window as the duration of the spoof (i.e., identical to the 'during' part); (2) 10 seconds; (3) 30 seconds; (4) 1 minute and (5) 5 minutes. Normality is assumed under the central limit theorem. Levene's test indicated variances are not equal, resulting in the use of Welch's *t* tests to measure if liquidity was significantly different before, during and after the spoof for all five time windows. APMs for the *t* tests are calculated per 10-ms snapshot.

4 | RESULTS

First, a single JPM spoofing case is used to showcase the benefits of using message-based visualizations rather than time-based visualizations. Subsequently, one JPM spoofing case is discussed per spoofing category. The spoofing actions as identified by the CFTC (CFTC, 2020) are described in detail, followed by LOB visualizations of these actions in subsections to facilitate the reading and interpretation of the figures. We examine four dimensions—trades, volume, cancellations, and liquidity—to show how they behave during a spoof.

4.1 | Snapshots versus message-based visualizations

The spoofing of JPM in the September 2015 Ultra T-Bond is used to illustrate the benefits of message-based visualizations. In summary, this spoofing involved one iceberg genuine order with one contract displayed and 199 contracts hidden on the bid side and a single ask spoof order of 100 contracts. More details of this particular spoof are discussed in Section 4.3. The contract is visualized during the spoofing time window on 30 June 2015 between 08:45:40 and 08:46:10.⁶ Rather than using the quoted five-decimal prices, prices in the visualizations are

⁶Interactive data visualizations can be included for each figure to let readers interact and engage with our research. Code can also be made available to readers.




FIGURE 1 Visualization of Ultra T-Bond September 2015 limit order book (LOB) using 5-s snapshots. The top panel shows the volumes at the individual bid and ask levels between prices of 153.5 and 154 points. Each unit on the x-axis is 1 s. The y-axis represents the price of the Ultra T-Bond in points. The colour represents the volume at each price level of the LOB for each 5-s snapshot. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The bottom panel shows the cumulative trade volume per second. A steeper (flatter) line signals a higher (lower) rate of traded volume. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the spoof order of 100 contracts was placed, when the first contract of the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

rounded to two decimals for readability. Figure 1 shows the behaviour of the Ultra T-Bond LOB using 5-s snapshots. The top panel shows the 10 ask (bid) levels above (below) the midpoint price, as indicated by the red horizontal line. The colours show the volumes at each price level. The various spoofing actions as identified by the CFTC (CFTC, 2020) are marked by vertical red lines. The bottom panel visualizes cumulative trade volume. Figure 1 demonstrates that the spoofing remains invisible when using high-frequency data and visualizing it using 5-s snapshots. The volume remains relatively constant at the individual bid and ask levels and the midpoint price is also relatively constant. The placing and cancelling of the spoof order happened within the same snapshot interval, leaving the addition and subtraction of 100 contracts invisible. Cumulative trade volume increases in a staircase pattern at the end of every 5-s snapshot. The only visible spoofing-related action in Figure 1 is the significant increase in trading volume 25 seconds into the time window, that is, the 51 contracts from the genuine order that were executed. However, one would not know that this was spoofing from merely looking at this figure.

Figure 2 is identical to Figure 1 but uses 1-s snapshots instead of 5-s snapshots. Contrary to Figure 1, the spoofing activities are visible in Figure 2. The addition and cancellation of the spoof order are now visible as a yellow bar at the first ask level, whereas they were not in Figure 1. Furthermore, trading volume spikes when the genuine order is executed and increases more gradually. However, this visualization cannot convey the exact timing of the spoofing activities. For




FIGURE 2 Visualization of Ultra T-Bond September 2015 limit order book (LOB) using 1-s snapshots. The top panel shows the volumes at the individual bid and ask levels between prices of 153.5 and 154 points. Each unit on the x-axis is 1 s. The y-axis represents the price of the Ultra T-Bond in points. The colour represents the volume at each price level of the LOB for each 1-s snapshot. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The bottom panel shows the cumulative trade volume per second. A steeper (flatter) line signals a higher (lower) rate of traded volume. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the spoof order of 100 contracts was placed, when the first contract of the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

example, the spoof order was cancelled at 08:46:04.418, but the visualization's 1-s resolution shows it as having been cancelled 'some time between 08:46:04 and 08:46:05'. Due to this lower resolution, the spoofing order appears to have been active for a longer time period than it actually was, as shown by the yellow bar after the vertical red line that reads 'Spoof Order cancelled'.

Visualizing the LOB using 1-ms snapshots would solve the problem of the delay between trading action and visualization, as the granularity of the visualization equals that of the timestamps in the raw data (i.e., 1 ms). However, the exchange frequently receives multiple messages, that is, changes to the LOB, within the same millisecond. Hence, information may be lost, as changes within the same millisecond are aggregated and not individually visible. Therefore, Figure 3 visualizes the LOB using messages instead of time-based snapshots. An additional panel is added to the bottom of Figure 3 to indicate how time passes between messages (blue line) and when one second has passed (green line). A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The LOB volume shows more information on (small) volume changes than the figures before. The addition and subtraction of small volumes might be an indication of algorithms 'probing' for other algorithms and hidden liquidity (Bongiovanni et al., 2006; Chakrabarty & Shaw, 2008). These 'probes' become visible when using (almost) the same data granularity as trading algorithms, that is, using messages rather than millisecond snapshots. Visualizations based on messages show what is actually happening in the market. In addition, they make the effect of executing an iceberg order more visible. The JPM spoofing in the Ultra

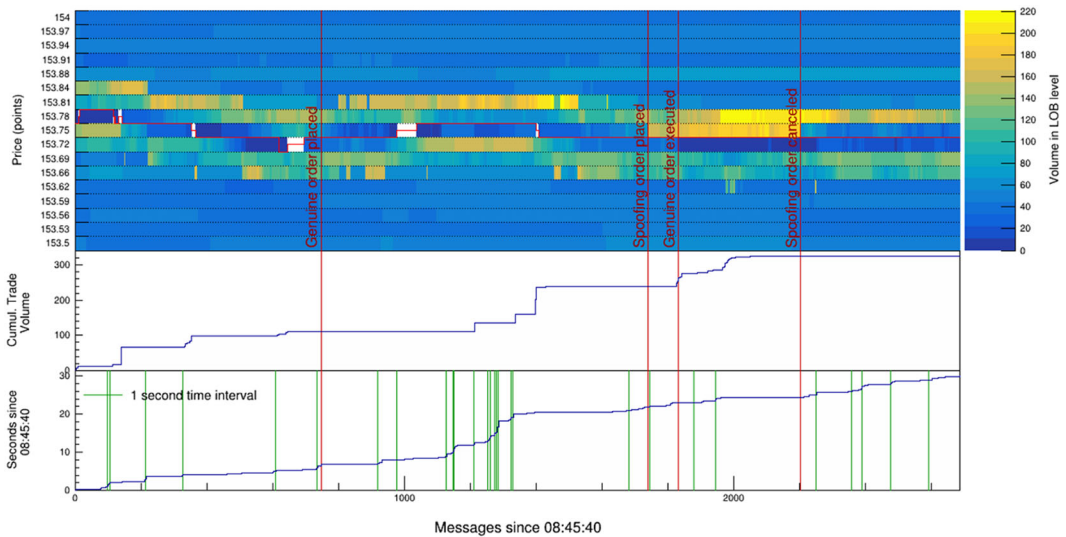


FIGURE 3 Visualization of Ultra T-Bond September 2015 limit order book (LOB) using messages received by the exchange. The top panel shows the volumes at the individual bid and ask levels between prices of 153.5 and 154 points. Each unit on the x-axis is one message. The y-axis represents the price of the Ultra T-Bond in points. The colour represents the volume at each price level of the LOB for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The middle panel shows the cumulative trade volume per message. A steeper (flatter) line signals a higher (lower) rate of traded volume. The bottom panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The green vertical lines indicate when 1 s has passed. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the spoof order of 100 contracts was placed, when the first contract of the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

T-Bond market consisted of a genuine iceberg order, and this becomes visible in the cumulative trade volume panel, once the first contract of the genuine order is executed. Many trades take place within the same millisecond, which would be aggregated (into one trade) in a snapshot-based visualization. However, Figure 3 shows that trade volume accumulated slower in this event, as the iceberg order was executed one contract at a time. This information was not visible in the previous visualizations and can help to understand spoofing behaviour.

4.2 | Traditional spoofing

Two futures contracts are part of the ‘traditional spoofing’ category: the March 2010 and December 2011 contracts from the 10-Year T-Note market. This section only discusses and presents results for the December 2011 contract, as both contracts show similar results. Table 1 shows the spoofing actions of the December 2011 contract, which took a total of 3.749 seconds. The spoof consisted of the placement of one genuine order on the first level of the bid side and a single large spoof order on the first ask level.

TABLE 1 Spoofing actions on 27 September 2011 in the 10-Year T-Note December 2011 futures market

This table reports the various spoofing actions JPMorgan took on 27 September 2011 in the 10-Year T-Note December 2011 futures market. Per spoof action, the table reports the timestamp (*Time*), whether it concerned a genuine or spoof order (*Order Type*), the limit order book (LOB) side the spoof action occurred on (*LOB Side*), whether the order from the spoof action was added or cancelled (*Action*), the price level affected by the spoof action (*Price (points)*) and the volume related to the spoof action (*Volume*).

Time	Order type	LOB side	Action	Price (points)	Volume
14:03:54.205	Genuine order	Bid	Add	129.578125	50
14:03:57.636	Spoof order	Ask	Add	129.59375	3000
14:03:57.671	<i>Complete genuine order executed</i>				
14:03:57.954	Spoof order	Ask	Cancel	129.59375	3000

TABLE 2 LOB state 1 ms before placement of the genuine order from the 10-Year T-Note December 2011 spoof

This table reports the state of the limit order book (LOB) 1-ms before the genuine order from the 10-Year T-Note December 2011 spoof was added. It shows the prices and volumes of each level on the bid and ask side.

Bid volume	Bid price (points)	Level	Ask price (points)	Ask volume
431	129.578	1	129.594	640
1889	129.562	2	129.609	1415
1742	129.547	3	129.625	1593
1720	129.531	4	129.641	1201
1648	129.516	5	129.641	1201
1893	129.5	6	129.641	1201
1041	129.484	7	129.641	1201
979	129.469	8	129.703	953
592	129.453	9	129.719	699
1081	129.438	10	129.734	658

Table 2 shows the state of the LOB one millisecond before the first spoofing action, providing insight into what would have happened if JPM had placed the genuine order as a market order rather than a limit order. The spoofing involved buying 50 contracts at 129.578125 points, with a total underlying value of \$6,478,906.25 (one point equalling \$1000). Had the same number of contracts been bought through a market order, JPM would have bought at 129.594 points, representing a total underlying value of \$6,479,700. Excluding trading costs, JPM thus succeeded in buying the contracts \$793.75 cheaper through spoofing than without spoofing.

Notably, in both traditional spoofing cases, the genuine order was placed on the first level of the bid side. Hence, these spoofing actions might not have been used to move the price—as otherwise

the genuine order would have been placed on a deeper level of the LOB⁷—but to attract more liquidity to the market to sell at the price of the first bid level. At the time the genuine order was placed, the first bid level already comprised 431 contracts. Hence, due to the price–time–priority rule, 431 contracts had to be sold at 129.578 points first, before the 50 contracts of the genuine order could be sold. This hypothesis—the motivation of this spoof being to attract liquidity—is further examined in Sections 4.2.4 and 4.6.

4.2.1 | Traditional spoofing: Visualization of the LOB and trades around spoofing

Figure 4 shows the behaviour of the LOB and trades around the spoofing of the December 2011 contract between 13:03:45 and 13:04:05. The top panel shows the last traded price (blue line) and the occurrence of individual trades (grey lines). The second and third panel visualize the LOB and cumulative trade volume, respectively. The bottom panel shows the number of messages reported by the exchange in the relevant time window and, hence, the amount of time that passes between messages.

The second panel in Figure 4 shows that, when the genuine order was added, individual LOB levels contained volumes of between 500 and 2500 contracts.⁸ When the spoof order of 3000 contracts was placed, volume on the first ask level increased significantly, as indicated by the bright yellow colour. This increase in volume remained in the LOB during the execution of the genuine order and ended when the spoof order was cancelled. The addition of the spoof order, the execution of the genuine order and the cancellation of the spoof order all occurred within the same second, as indicated by the space between the green vertical lines.

The top panel in Figure 4 shows that when the genuine bid order was placed at 129.578 points, the last traded price was also 129.578 points. This illustrates once more that the goal of this spoof may not have been to move the price, but to attract more liquidity, so as to increase the chance of fully executing the genuine bid order of 50 contracts.⁹ This will be further explored in Sections 4.2.4 and 4.6. The last traded price remained constant at 129.578 points during all spoofing actions. The cumulative trade volume panel in Figure 4 shows that no trades took place in the time window until the genuine order and spoof order were placed.¹⁰ After the spoof order was placed, a staircase pattern emerged. Our data shows that this was caused by the genuine bid order not being executed at once but being split into smaller executed trades. After the genuine order was fully executed, cumulative trade volume continued to increase—albeit at a lower volume—and remained constant (i.e., no trades occurred) right before and after the cancellation of the spoof order.

⁷In this event, spoofing would be used to move the price in the desired direction and push it through the first level(s) of the LOB to get a better price than the current best bid/ask.

⁸Volume in The March 2010 LOB was considerably higher, as most levels contained volumes of between 1500 and 3500 contracts.

⁹In contrast to the March 2010 contract, where the last traded price (118.281 points) was higher when the genuine order was added (118.266 points).

¹⁰This does not mean no trades occurred in the market during that day, but that no trades occurred in the visualized time window until the spoof order was placed.

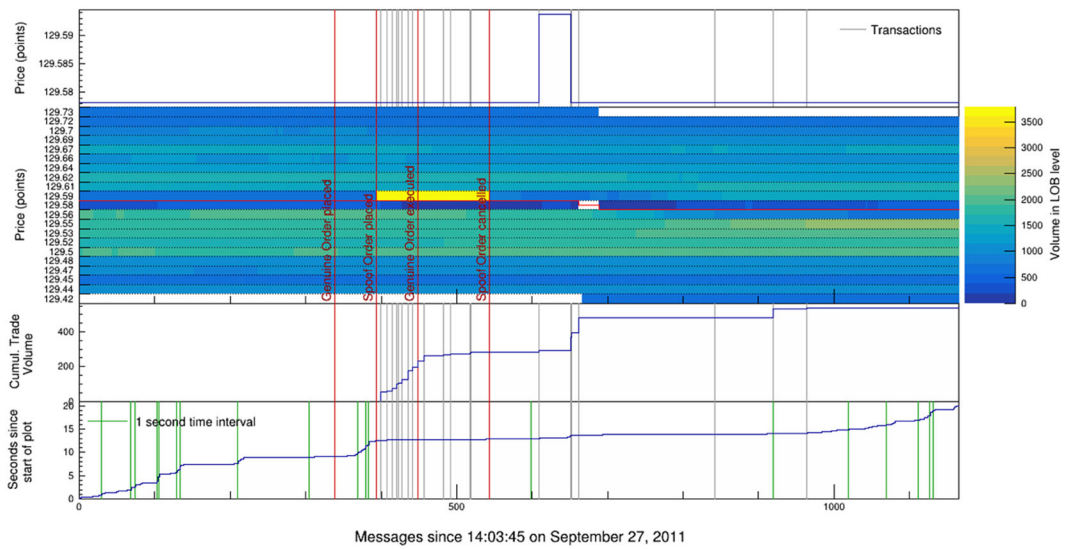


FIGURE 4 Visualization of the limit order book (LOB) and trade behaviour around the spoof of 27 September 2011 in the 10-Year T-Note December 2011 futures market. The first panel shows the price of the last trade that took place (blue line) and when a trade took place (grey line). The second panel shows the volumes at the individual bid and ask levels between prices of 129.42 and 129.73 points. Each unit on the x-axis is one message. The y-axis represents the price of the 10-Year T-Note in points. The colour represents the volume at each price level of the LOB for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The third panel shows the cumulative trade volume per second. The fourth panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the spoof order of 3000 contracts was placed, when the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

4.2.2 | Traditional spoofing: Visualization of volume around spoofing

Figure 5 visualizes the volume changes on the first bid and ask levels around the time of the spoof. When the genuine order was added, volume on the first bid and ask levels was stable at approximately 480 and 640 contracts, respectively. Volume increased significantly by 3000 contracts on the first ask level when the spoof order was added. Between the spoof order being added and the genuine order being executed, the volume on the first bid level decreased gradually. This is attributed to trades being executed and taking volume from the bid level, as shown in the third panel in Figure 4. After the genuine order was executed, volume on the first bid level decreased to two contracts. When the spoof order was cancelled, volume on the first ask level decreased significantly by 3000 contracts to 771 contracts and volume on the first bid level gradually increased.




FIGURE 5 Visualization of the first-level bid and ask volume behaviour around the spoof of 27 September 2011 in the 10-Year T-Note December 2011 futures market. The first panel shows the volume of the best ask level. The second panel shows the volumes at the individual bid and ask levels between prices of 129.42 and 129.73 points. Each unit on the x-axis is one message. The y-axis represents the price of the 10-Year T-Note in points. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The third panel shows the volume of the best bid level. The fourth panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the spoof order of 3000 contracts was placed, when the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

4.2.3 | Traditional spoofing: Visualization of cancellations around spoofing

Figure 6 visualizes the cancellations on the first ask (top panel) and bid levels (third panel) around the time of the spoof. Cancellations on the first ask level remained close to zero until the cancellation of the spoof order, only to increase significantly after the spoof order of 3000 contracts was removed from the LOB. Cancellations on the first bid level remained constant at a cumulative cancellation volume of around 300 contracts during all spoofing actions.

4.2.4 | Traditional spoofing: Visualization of liquidity around spoofing

The first and third panels in Figure 7 show the behaviour of liquidity costs on the ask and bid side, respectively, around the December 2011 spoof. On the ask side, liquidity costs were relatively stable at around 4.7 bps up until the spoof order was placed. When the spoof order was added, liquidity costs drastically decreased to approximately 2.2 bps. After the cancellation of the spoof order, liquidity costs returned to approximately the same level as before the spoof




FIGURE 6 Visualization of cumulative first-level bid and ask cancellation volume around the spoof of 27 September 2011 in the 10-Year T-Note December 2011 futures market. The first panel shows the cumulative volume of cancellations of the best ask level. The second panel shows the volumes at the individual bid and ask levels between prices of 129.42 and 129.73 points. Each unit on the x -axis is one message. The y -axis represents the price of the 10-Year T-Note in points. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The third panel shows the cumulative volume of cancellations of the best bid level. The fourth panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the spoof order of 3000 contracts was placed, when the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

order was placed. Bid-side liquidity costs remained relatively stable between 3.9 and 4.2 bps during all spoofing actions by JPM.¹¹

Welch's t tests were used to test whether liquidity differed significantly between the periods before, during and after the spoofing. Results are reported in Table 3. In any time window, liquidity costs before the spoofing were higher than during the spoofing. In other words, liquidity was lower before than during the spoofing and improved during the spoof. After the spoof ended, liquidity costs significantly increased and, hence, liquidity was significantly lower after than during the spoof. Up to 30 seconds after the spoof ended, liquidity costs were higher than before the spoof. In other words, liquidity was significantly worse after the spoof than before.¹²

¹¹The March 2010 contract shows more fluctuations in liquidity on the bid and ask sides than the December 2011 contract, as other volume not related to the JPM spoofing example was repeatedly shifted between the 10th bid and 10th ask level.

¹²Results differ for the March 2010 contract, as can be seen in the Online Supporting Information Appendix.




FIGURE 7 Visualization of bid and ask liquidity costs (Adverse Price Movement [APM]) behaviour around the spoof of 27 September 2011 in the 10-Year T-Note December 2011 futures market. The first panel shows the APM of the ask side. APM measures the liquidity costs (in basis points) of a trader who wants to buy or sell a specific dollar value by submitting market orders. The second panel shows the volumes at the individual bid and ask levels between prices of 129.42 and 129.73 points. Each unit on the x-axis is one message. The y-axis represents the price of the 10-Year T-Note in points. The colour represents the volume in each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The third panel shows the APM for the bid side. The fourth panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the spoof order of 3000 contracts was placed, when the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 3 Mean ask liquidity costs (bps) around the 10-Year T-Note December 2011 spoof for different time windows

This table reports the mean liquidity costs (basis point [bps]), measured by Adverse Price Movement [APM] around the spoof in the 10-Year T-Note December 2011 market for different periods and various time windows. *Before* represents the time up until the spoof order was added; *During* the period from when the spoof order was added until it was cancelled and *After* the time following the cancellation of the spoof order. Five different time windows are used, the *Spoof Duration* time window being 0.310 seconds. A lower APM indicates that liquidity costs are low and, hence, liquidity is high. Welch's *t* tests were used to test for mean differences between the periods. Significance at the 0.1% and 5% (two-tailed) levels is indicated by *** and *, respectively.

Time window	Before versus during	During versus after	Before versus after
Spoof duration	4.587 > 2.250***	2.250 < 4.620***	4.587 < 4.620***
10 seconds	4.605 > 2.250***	2.250 < 4.838***	4.605 < 4.838***
30 seconds	4.623 > 2.250***	2.250 < 4.632***	4.623 < 4.632*
1 minute	4.700 > 2.250***	2.250 < 4.680***	4.700 > 4.680***
5 minutes	4.644 > 2.250***	2.250 < 4.945***	4.644 < 4.945***

4.3 | Traditional spoofing with iceberg orders

Two futures contracts are part of the ‘traditional spoofing with iceberg orders’ category: the Silver March 2014 and Ultra T-Bond September 2015 contracts. This section only discusses results for the Ultra T-Bond September 2015 contract. Table 4 outlines the spoofing actions JPM undertook in the Ultra T-Bond September 2015 market on 30 June 2015. All spoofing actions lasted for 21.447 seconds and consisted of a single genuine and spoof order. The genuine order was an iceberg order on the first bid level and consisted of one visible contract and 199 hidden contracts. The spoof order was placed on the first ask level and consisted of 100 contracts.

Table 5 shows the state of the LOB one millisecond before JPM's first spoofing action in the Ultra T-Bond market. The spoofing involved buying 51 contracts at 153.71875 points, representing a total underlying value of \$7,839,656.25 (one point equalling \$1000). If the JPM trader had executed their genuine order with market orders, they would have bought 51 contracts at 153.75 points, representing a total underlying value of \$7,841,250. Hence, due to spoofing, JPM bought the contracts \$1593.75 cheaper, excluding trading costs. Assuming the JPM trader wanted the full genuine order executed, that is, buy 200 contracts rather than 51 contracts, the gains would have been larger. In that situation, a market order of 200 contracts would have ‘run up’ the LOB: they would have bought 69 contracts at 153.75 points; 127 contracts at 153.78125 points and four contracts at 153.8125 points. The total underlying value using market orders would have been \$30,754,218.75, which is \$10,468.75 more than the total underlying value of buying 200 contracts in the spoofing scenario (\$30,743,750). JPM might have placed an iceberg order or initiated the spoofing actions not to move the price, but to attract more liquidity to avoid running up the LOB and incur liquidity costs. This will be further explored in Sections 4.3.4 and 4.6.

TABLE 4 Spoofing actions on 30 June 2015 in the Ultra T-Bond September 2015 futures market

This table reports the various spoofing actions JPMorgan took on 30 June 2015 in the Ultra T-Bond September 2015 futures market. Per spoof action, the table reports the timestamp (*Time*), whether it concerned a genuine or spoof order (*Order Type*), the limit order book (LOB) side the spoof action occurred on (*LOB Side*), whether the order from the spoof action was added or cancelled (*Action*), the price level affected by the spoof action (*Price (points)*) and the volume related to the spoof action (*Volume*).

Time	Order type	LOB side	Action	Price (points)	Volume
08:45:46.627	Genuine order	Bid	Add	153.71875	1 displayed 199 hidden
08:46:01.891	Spoof order	Ask	Add	153.75	100
08:46:02.979	<i>First contract of genuine order executed</i>				
08:46:04.288	<i>Last contract of genuine order executed (51 of 200 contracts executed)</i>				
08:46:04.418	Spoof order	Ask	Cancel	153.75	100
08:46:08.074	Genuine order	Bid	Cancel	153.71875	149

TABLE 5 LOB state 1 ms before placement of the genuine order from the Ultra T-Bond September 2015 spoof

This table reports the state of the limit order book (LOB) 1 ms before the genuine order from the Ultra T-Bond September 2015 spoof was added. It shows the prices and volumes of each level on the bid and ask side.

Bid volume	Bid price (points)	Level	Ask price (points)	Ask volume
30	153.71875	1	153.75000	69
121	153.68750	2	153.78125	127
104	153.65625	3	153.81250	52
37	153.62500	4	153.84375	40
45	153.59375	5	153.87500	65
42	153.56250	6	153.90625	43
41	153.53125	7	153.93750	45
54	153.50000	8	153.96875	51
47	153.46875	9	154.00000	38
47	153.43750	10	154.03125	46

4.3.1 | Traditional spoofing with iceberg orders: Visualization of the LOB and trades around spoofing

Figure 8 visualizes the behaviour of the LOB and information about trades around the spoofing in the Ultra T-Bond September 2015 contract, on 30 June 2015 between 08:45:40 and 08:46:10. When the genuine order was added, most of the volume in the LOB was concentrated on the second and third bid levels and the first two ask levels. Once the spoof order of 100 contracts was placed, volume increased significantly on the first ask level, as indicated by a bright yellow colour. After the genuine order was executed, volume on the first bid level decreased, indicated by ever darker shades of blue. In contrast, more volume was added on the second ask level. Volume on the first ask level was significantly lower once the spoof order was cancelled.

The top panel in Figure 8 shows that the price of the genuine order and the last traded price were identical (153.71875 points) at the time of placing the genuine order. Hence, the spoof order may have been used to attract more liquidity to the price of the genuine order, which will be further explored in Sections 4.3.4 and 4.6. When the spoof order was placed, the last traded price was 153.75 points, and shortly after the placement—1.087 seconds later—it decreased to the price level of the genuine order, to stay there for the remainder of the visualized time window. The cumulative trade volume panel in Figure 8 provides more information about the trading patterns of iceberg orders: while previous trades showed staircase patterns, the spoofing-related trades are more gradual because of the associated iceberg order. This order only executes one trade at a time, whereby each trade is recorded in a separate message.¹³ Hence, in this case, visualizing trades based on messages provides more insight into the type of

¹³The iceberg order of the Silver March 2014 spoof used five visible contracts and, hence, five contracts at a time can be executed. This caused cumulative volume to increase in a staircase pattern rather than gradually, as it did in the Ultra T-Bond September 2015 contract.




FIGURE 8 Visualization of the limit order book (LOB) and trade behaviour around the spoof of 30 June 2015 in the Ultra T-Bond September 2015 futures market. The *first* panel shows the price of the last trade that took place (blue line) and when a trade took place (grey line). The *second* panel shows the volumes at the individual bid and ask levels between prices of 153.5 and 154 points. Each unit on the x-axis is one message. The y-axis represents the price of the Ultra T-Bond in points. The colour represents the volume at each price level of the LOB for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the cumulative trade volume per second. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the spoof order of 100 contracts was placed, when the first contract of the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

order and trade. Furthermore, it can provide additional insights into the type of trader. For example, an algorithm could also have produced the same type of trade pattern, as algorithms trade in nanoseconds and can, therefore, rapidly execute market orders in a short time window.

4.3.2 | Traditional spoofing with iceberg orders: Visualization of volume around spoofing

Figure 9 visualizes the changes in volume on the first bid and ask levels around the spoofing in the Ultra T-Bond September 2015 contract. When the genuine order was added, volume on the first bid and ask levels changed regularly, which can be attributed to a new first price level being added or removed from the LOB. When the spoof order was added, volume on the first ask level increased by 100 contracts and kept increasing gradually until the spoof order was removed. Volume on the first bid level remained relatively stable when the spoof order was added and dropped when the genuine order was executed. At this point, it remained between one and 10 contracts until the spoof order was cancelled and shortly after.




FIGURE 9 Visualization of first-level bid and ask volume behaviour around the spoof of 30 June 2015 in the Ultra T-Bond September 2015 futures market. The *first* panel shows the volume of the best ask level. The *second* panel shows the volumes at the individual bid and ask levels between prices of 153.5 and 154 points. Each unit on the x-axis is one message. The y-axis represents the price of the Ultra T-Bond in points. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the volume of the best bid level. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the spoof order of 100 contracts was placed, when the first contract of the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

4.3.3 | Traditional spoofing with iceberg orders: Visualization of cancellations around spoofing

Figure 10 shows the cumulative cancellations on the first bid and ask levels around the spoof. In general, the volume cancelled on the first bid and ask levels is small. Up until when the spoof order was cancelled, cancellations on the first ask level were increasing gradually. When the spoof order was cancelled, it increased significantly by 100 contracts. Cancellations on the first bid level continued to gradually increase in the visualized time window. Figure 10 complements Figure 9, in that Figure 10 explains whether the shifts in Figure 9 should be attributed to cancellations or to other causes.

4.3.4 | Traditional spoofing with iceberg orders: Visualization of liquidity around spoofing

Figure 11 visualizes the bid and ask liquidity costs around the spoof in the Ultra T-Bond September 2015 contract. Before the spoof order was placed, liquidity costs on the ask side fluctuated between 9.5 and 13 bps. Immediately when the spoof order was placed, ask liquidity costs dropped from




FIGURE 10 Visualization of cumulative first-level bid and ask cancellation volume around the spoof of 30 June 2015 in the Ultra T-Bond September 2015 futures market. The *first* panel shows the cumulative volume of cancellations of the best ask level. The *second* panel shows the volumes at the individual bid and ask levels between prices of 153.5 and 154 points. Each unit on the x-axis is one message. The y-axis represents the price of the Ultra T-Bond in points. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the cumulative volume of cancellations of the best bid level. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the spoof order of 100 contracts was placed, when the first contract of the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

10.38 to 7.97 bps and further decreased to approximately 6 bps right before the spoof order was cancelled. After the spoof order was cancelled, ask liquidity costs fluctuated between 9 and 12 bps in the visualized time window. Compared to the ask side, bid side liquidity costs were relatively more volatile, fluctuating between 9 and 13 bps.

Table 6 shows the test results for whether liquidity costs were significantly different before, during and after the spoofing. Irrespective of the time window, liquidity costs were higher before and after the spoof than during the spoof. In other words, liquidity was better during the spoof than before and after. When comparing liquidity costs before and after the spoof, the results differ per time window. Liquidity was better 2.52 seconds after the spoof than before the spoof. For each subsequent time window, the results are mixed.

4.4 | Layered spoofing

Four futures contracts are part of the 'layered spoofing' category: the Silver March 2012, Silver May 2014, Gold April 2014 and T-Bond September 2009 contracts. This section only discusses results for




FIGURE 11 Visualization of bid and ask liquidity costs (Adverse Price Movement [APM]) behaviour around the spoof of 30 June 2015 in the Ultra T-Bond September 2015 futures market. The *first* panel shows the APM of the ask side. APM measures the liquidity costs (in basis points) of a trader who wants to buy or sell a specific dollar value by submitting market orders. The *second* panel shows the volumes at the individual bid and ask levels between prices of 153.5 and 154 points. Each unit on the *x*-axis is one message. The *y*-axis represents the price of the Ultra T-Bond in points. The colour represents the volume in each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the APM for the bid side. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the spoof order of 100 contracts was placed, when the first contract of the genuine order was executed and when the spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

the T-Bond September 2009 contract. Table 7 shows the spoofing actions by JPM in the T-Bond September 2009 contract (CFTC, 2020), which lasted for a total of 8.706 seconds. The spoof consisted of one genuine order with a volume of 100 contracts at the second ask level¹⁴ and six spoof orders with a volume of 300 contracts.

Table 8 shows the state of the T-Bond September 2009 contract on 20 July 2009 1 ms before the genuine order was placed. JPM sold 100 contracts at 116.171875 points, amounting to a total underlying value of \$11,617,187.5. Had JPM submitted their genuine order as a market order rather than a limit order, it would have consumed the first and part of the second bid level. In that scenario, JPM would have sold 59 contracts at 116.141 points and 41 contracts at 116.125 points, representing a total underlying value of \$11,613,444. Hence, JPM sold their contracts for \$3743.5 more through spoofing, excluding trading costs.

¹⁴The genuine orders for the Silver March 2012, Silver May 2014 and Gold April 2014 contracts were all placed on the first rather than the second ask level.

TABLE 6 Mean ask liquidity costs (bps) around Ultra T-Bond September 2015 spoof for different time windows

This table reports the mean liquidity costs (basis point [bps], measured by Adverse Price Movement [APM]) around the spoof in the Ultra T-Bond September 2015 market for different periods and various time windows. *Before* represents the time up until the spoof order was added; *During* the period from when the spoof order was added until it was cancelled and *After* the time following the cancellation of the spoof order. Five different time windows are used, the *Spoof Duration* time window being 2.52 seconds. A lower APM indicates that liquidity costs are low and, hence, liquidity is high. Welch's *t* tests were used to test for mean differences between the periods. Significance at the 0.1% (two-tailed) level is indicated by ***.

Time window	Before versus during	During versus after	Before versus after
Spoof duration	10.695 > 7.981***	7.981 < 10.267***	10.695 > 10.267***
10 seconds	10.545 > 7.981***	7.981 < 10.851***	10.545 < 10.851***
30 seconds	10.873 > 7.981***	7.981 < 10.907***	10.873 = 10.907
1 minute	10.935 > 7.981***	7.981 < 10.799***	10.935 > 10.799***
5 minutes	13.284 > 7.981***	7.981 < 10.369***	13.284 < 10.369***

TABLE 7 Spoofing actions on 20 July 2009 in the T-Bond September 2009 futures market

This table reports the various spoofing actions JPMorgan took on 20 July 2009 in the T-Bond September 2009 futures market. Per spoof action, the table reports the timestamp (*Time*), whether it concerned a genuine or spoof order (*Order Type*), the limit order book (LOB) side the spoof action occurred on (*LOB Side*), whether the order from the spoof action was added or cancelled (*Action*), the price level affected by the spoof action (*Price (points)*) and the volume related to the spoof action (*Volume*).

Time	Order type	LOB side	Action	Price (points)	Volume
07:47:13.597	Genuine order	Ask	Add	116.171875	100
07:47:17.098	Spoof layer 1	Bid	Add	116.078	300
07:47:17.847	Spoof layer 2	Bid	Add	116.094	300
07:47:18.583	Spoof layer 3	Bid	Add	116.109	300
07:47:19.379	Spoof layer 4	Bid	Add	116.125	300
07:47:20.212	Spoof layer 5	Bid	Add	116.141	300
07:47:21.020	Spoof layer 6	Bid	Add	116.156	300
07:47:21.036	<i>Complete genuine order executed</i>				
07:47:22.039	Spoof layer 6	Bid	Cancel	116.156	300
07:47:22.064	Spoof layer 5	Bid	Cancel	116.141	300
07:47:22.064	Spoof layer 4	Bid	Cancel	116.125	300
07:47:22.064	Spoof layer 3	Bid	Cancel	116.109	300
07:47:22.067	Spoof layer 2	Bid	Cancel	116.094	300
07:47:22.303	Spoof layer 1	Bid	Cancel	116.078	300

TABLE 8 LOB state 1 ms before placement of the genuine order from the T-Bond September 2009 spoof

This table reports the state of the limit order book (LOB) 1 ms before the genuine order from the T-Bond September 2009 spoof was added. It shows the prices and volumes of each level on the bid and ask side.

Bid volume	Bid price (points)	Level	Ask price (points)	Ask volume
59	116.141	1	116.156	55
85	116.125	2	116.172	62
90	116.109	3	116.188	180
163	116.094	4	116.203	102
79	163.078	5	116.219	105
116	116.062	6	116.234	108
75	116.047	7	116.25	61
184	116.031	8	116.266	204
42	116.016	9	116.281	233
35	116	10	116.297	41

4.4.1 | Layered spoofing: Visualization of the LOB and trades around spoofing

Figure 12 shows the visualization of the LOB and trades around the JPM spoofing in the T-Bond September 2009 market on 20 July 2009 from 07:47:10 to 07:47:30. The second panel shows that when the genuine order was added, individual levels contained approximately between 50 and 250 contracts. Most volume was concentrated on the third, eighth and ninth ask levels and on the fourth and eighth bid levels. Spoof orders of 300 contracts were placed on six different levels and, as indicated by the bright yellow colour, were relatively large compared to the volumes on these levels. The spoof orders were placed from the lower to the higher levels in the LOB, that is, from level six to level one. Conversely, spoof orders were cancelled from the higher to the lower levels in the LOB, that is, from level one to level six. Hence, the spoof orders closest to the top of the LOB were active for the shortest amount of time. The execution of the genuine order and the cancellation of all spoof orders occurred within the same second, as indicated by the green vertical lines in the lower panel.

The top panel in Figure 12 shows that, when the genuine order was placed at 116.171875 points (rounded 116.172 points), the last traded price was 116.141 points. Hence, the goal of this spoof might have been to move the price up towards the ask price of the genuine order.¹⁵ Before the first spoof order was placed, the last traded price moved between the highest bid (116.141 points) and lowest ask (116.156 points).¹⁶ This illustrates which side triggers the trade: a trader wanting to buy and taking the lowest ask, or a trader wanting to sell and taking the highest bid. Shortly after the fifth spoof order was placed, the trade price increased to 116.172 points and the

¹⁵The price of the genuine order was equal to the last traded price in the case of the Silver May 2014 spoof.

¹⁶The last traded price of the Silver May 2014 contract did not move during the visualized time window (from 08:18:35 to 08:18:50).




FIGURE 12 Visualization of the limit order book (LOB) and trade behaviour around the spoof of 20 July 2009 in the T-Bond September 2009 futures market. The *first* panel shows the price of the last trade that took place (blue line) and when a trade took place (grey line). The *second* panel shows the volumes at the individual bid and ask levels between prices of 116 and 116.33 points. Each unit on the x-axis is one message. The y-axis represents the price of the T-Bond in points. The colour represents the volume at each price level of the LOB for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the cumulative trade volume per second. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the first spoof order of 300 contracts was placed, when the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

genuine order was executed. The cumulative trade panel in Figure 12 shows that, before the genuine order was executed, one large trade occurred (shortly after the genuine order was added) while the other trades were relatively small. At the time of the execution of the genuine order and after, larger trades were executed, as indicated by the staircase pattern.

4.4.2 | Layered spoofing: Visualization of volume around spoofing

Figure 13 visualizes the changes in volume on the second levels around the spoof of the T-Bond September 2009 contract. When the genuine order was added, the second ask level consisted of 62 contracts and the second bid level of 85 contracts. Both volumes remained relatively constant within these price levels until the first spoof order was added. Large fluctuations in the second ask level were mainly attributable to a changing bid-ask spread and, hence, changing second ask price level. Once the spoof order was placed on the second bid level, around the 380 message mark, the volume increased significantly by 300 contracts. Although the price level of the second bid level changed around the 480 and 500 message mark, the volume on the second bid level continued to be high as 300 contracts were added to

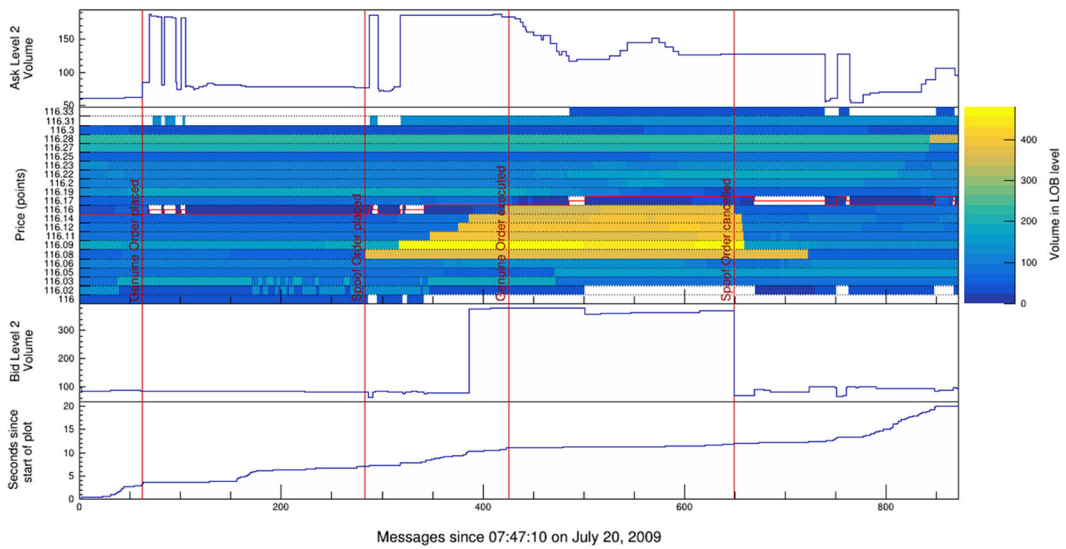


FIGURE 13 Visualization of second-level bid and ask volume behaviour around the spoof of 20 July 2009 in the T-Bond September 2009 futures market. The *first* panel shows the volume of the second ask level. The *second* panel shows the volumes at the individual bid and ask levels between prices of 116 and 116.33 points. Each unit on the x-axis is one message. The y-axis represents the price of the T-Bond in points. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the volume of the second bid level. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the first spoof order of 300 contracts was placed, when the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

multiple layers by the JPM trader. Once the spoof order was cancelled, the volume decreased significantly by 300 contracts.¹⁷

4.4.3 | Layered spoofing: Visualization of cancellations around spoofing

Figure 14 shows the cancellations on the second bid and ask levels around the spoof in the T-Bond September 2009 market. Cancellations on the second ask level gradually increased in the visualized time window, the largest cancellations being approximately 10 contracts in one message. Cumulative cancellations on the second bid level remained under 40 contracts up until the cancellation of the first spoof order. When the first spoof order was cancelled, it significantly increased by 300 contracts, after which it continued to gradually increase at a slower pace.

¹⁷Due to a frequently changing bid-ask spread in the visualized time window, the first bid and ask volumes fluctuated more in the Gold April 2014 contract than in the other spoofing examples.




FIGURE 14 Visualization of cumulative second-level bid and ask cancellation volume around the spoof of 20 July 2009 in the T-Bond September 2009 futures market. The *first* panel shows the cumulative volume of cancellations of the second ask level. The *second* panel shows the volumes at the individual bid and ask levels between prices of 116 and 116.33 points. Each unit on the x-axis is one message. The y-axis represents the price of the T-Bond in points. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the cumulative volume of cancellations of the second bid level. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the first spoof order of 300 contracts was placed, when the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

4.4.4 | Layered spoofing: Visualization of liquidity around spoofing

The ask and bid liquidity costs around the T-Bond September 2009 contract are visualized in Figure 15. The liquidity costs on the bid side fluctuated between 6.5 and 10 bps before the first spoof order was added and continued to decrease with every additional spoof order added, reaching their lowest point at 1.9 bps before stabilizing at approximately 3 bps. After all spoof orders were cancelled, the bid liquidity costs fluctuated between 4 and 9 bps. The ask liquidity costs fluctuated between 5.2 and 7.4 bps and reached their lowest point in the visualized time window during the spoof.¹⁸

Results from Welch's *t* tests for the T-Bond September 2009 spoof are reported in Table 9. The bid liquidity costs were significantly higher before and after the spoof than during the spoof, regardless of the time window. Hence, liquidity improved during the spoof. Up until 30 seconds after the spoof, the liquidity costs were significantly lower than during the spoof.

¹⁸The other spoofing examples in this category all showed a similar downward pattern in liquidity costs.




FIGURE 15 Visualization of bid and ask liquidity costs (Adverse Price Movement [APM]) behaviour around the spoof of 20 July 2009 in the T-Bond September 2009 futures market. The *first* panel shows the APM of the ask side. APM measures the liquidity costs (in basis points) of a trader who wants to buy or sell a specific dollar value by submitting market orders. The *second* panel shows the volumes at the individual bid and ask levels between prices of 116 and 116.33 points. Each unit on the x-axis is one message. The y-axis represents the price of the T-Bond in points. The colour represents the volume in each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the APM for the bid side. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine order was placed, when the first spoof order of 300 contracts was placed, when the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 9 Mean bid liquidity costs (bps) around the T-Bond September 2009 spoof for different time windows

This table reports the mean liquidity costs (basis point [bps]), measured by Adverse Price Movement [APM] around the spoof in the T-Bond September 2009 market for different periods and various time windows. *Before* represents the time up until the spoof order was added; *During* the period from when the spoof order was added until it was cancelled and *After* the time following the cancellation of the spoof order. Five different time windows are used, the *Spoof Duration* time window being 5.2 seconds. A lower APM indicates that liquidity costs are low and, hence, liquidity is high. Welch's *t* tests were used to test for mean differences between the periods. Significance at the 0.1% (two-tailed) level is indicated by ***.

Time window	Before versus during	During versus after	Before versus after
Spoof duration	7.788 > 4.284***	4.284 < 7.424***	7.788 > 7.424***
10 seconds	7.882 > 4.284***	4.284 < 7.340***	7.882 > 7.340***
30 seconds	6.723 > 4.284***	4.284 < 5.980***	6.723 > 5.980***
1 minute	5.954 > 4.284***	4.284 < 6.603***	5.954 < 6.603***
5 minutes	8.265 > 4.284***	4.284 < 7.791***	8.265 > 7.791***

TABLE 10 Spoofing actions on 22 June 2016 in the Platinum July 2016 futures market

This table presents the various spoofing actions JPMorgan took on 22 June 2016 in the Platinum July 2016 futures market. Per spoof action, the table reports the timestamp (*Time*), whether it concerned a genuine or spoof order (*Order Type*), the LOB side the spoof action occurred on (*LOB Side*), whether the order from the spoof action was added or cancelled (*Action*), the price level affected by the spoof action (*Price*) and the volume related to the spoof action (*Volume*).

Time	Order type	LOB side	Action	Price	Volume
02:14:33.935	Genuine order	Ask	Add	\$981.8	1 displayed 19 hidden
02:14:35.926	Spoof layer 1	Bid	Add	\$981.2	5
02:14:36.072	Spoof layer 2	Bid	Add	\$981.4	5
02:14:36.214	Spoof layer 3	Bid	Add	\$981.6	5
02:14:36.374	Spoof layer 4	Bid	Add	\$981.6	5
02:14:36.519	Spoof layer 5	Bid	Add	\$981.6	5
02:14:36.520	<i>Four contracts of genuine order executed</i>				
02:14:36.678	Spoof layer 6	Bid	Add	\$981.6	5
02:14:36.824	Spoof layer 7	Bid	Add	\$981.6	5
02:14:37.006	Spoof layer 8	Bid	Add	\$981.6	5
02:14:37.407	Spoof layer 3–8	Bid	Cancel	\$981.6	30
02:14:38.063	Spoof layer 2	Bid	Cancel	\$981.4	5
02:14:40.695	Spoof layer 1	Bid	Cancel	\$981.2	5

4.5 | Layered spoofing with iceberg orders

One futures contract is part of the ‘layered spoofing with iceberg orders’ category: the Platinum July 2016 contract. Table 10 outlines the spoofing actions JPM took in the Platinum market on 22 June 2016 (CFTC, 2020). All spoofing actions lasted for a total of 6.76 seconds and consisted of (1) a genuine iceberg order on the first ask level, with one contract displayed and nineteen hidden and (2) eight spoof orders with a volume of five contracts each.

Table 11 shows the state of the LOB 1 ms before adding the genuine iceberg order. JPM sold four contracts for \$981.80, with a total underlying value of \$196,360. Had they sold these four contracts with a market order, they would have sold two contracts for \$981.7 and two contracts for \$981.6, with a total underlying value of \$196,330. Hence, excluding trading costs, JPM received \$30 more by using a limit order and spoofing the market. Assuming that JPM wanted the full genuine iceberg order executed, that is, wanted to sell 20 rather than four contracts, the gains would have been larger. In that case, the spoofing would have resulted in JPM selling at an underlying value of \$981,800. Using a market order of volume 20, the order would have run down the LOB and consumed the first four bid levels. In that case, JPM would have sold at a total underlying value of \$981,400, which would have been \$400 less than with spoofing, excluding transaction costs.

TABLE 11 LOB state 1 ms before placement of the genuine order from the Platinum July 2016 spoof

This table reports the state of the limit order book (LOB) 1 ms before the genuine order from the Platinum July 2016 spoof was added. It shows the prices and volumes of each level on the bid and ask side.

Bid volume	Bid price	Level	Ask price	Ask volume
2	\$981.7	1	\$982.2	4
4	\$981.6	2	\$982.3	2
3	\$981.4	3	\$982.5	2
6	\$981.3	4	\$982.6	1
6	\$981.2	5	\$982.7	2
8	\$981.0	6	\$982.8	2
4	\$980.9	7	\$983.0	7
4	\$980.8	8	\$983.1	3
2	\$980.7	9	\$983.2	7
3	\$980.6	10	\$983.3	1

4.5.1 | Layered spoofing with iceberg orders: Visualization of the LOB and trades around spoofing

The behaviour of the LOB and trades around the spoofing of the Platinum July 2016 contract is visualized in Figure 16 between 02:14:25 and 02:14:45. When the genuine order was added on the first ask level, the volume on the individual LOB levels was low at between 0 and 10 contracts, as visualized in the second panel. The bid-ask spread was wider before the genuine order was added than after: \$0.4 and \$0.1, respectively. This may illustrate that the spoof orders were used by JPM to attract more liquidity to the market, thereby tightening the bid-ask spread. This will be further explored in Sections 4.5.4 and 4.6. The first spoof order was placed at the sixth bid level, the second spoof order at the fourth bid level and the third to eighth spoof orders at the second bid level. This is visualized in Figure 16 by a colour change on the respective level from blue to a lighter blue, green or yellow. Spoof orders were still being added 1 second after four contracts from the genuine order were executed, and the cancellations of the spoof orders started another second later.

The top panel in Figure 16 shows that, when the genuine order at price \$981.8 was added, a transaction occurred in the same millisecond at a trade price of \$981.8. Before this transaction, the last traded price was \$982.1. For the duration of JPM's spoofing actions, the transaction price remained at \$981.8. Cumulative trade volume increased steadily after the genuine order was placed.

4.5.2 | Layered spoofing with iceberg orders: Visualization of volume around spoofing

Figure 17 visualizes the volume changes in the second bid and ask levels around the time of the spoof. When the genuine order and the first spoof order were added, the volume on the second bid and ask level was low at two contracts on each side. Once spoof orders were added on the




FIGURE 16 Visualization of the limit order book (LOB) and trade behaviour around the spoof of 22 June 2016 in the Platinum July 2016 futures market. The *first* panel shows the price of the last trade that took place (blue line) and when a trade took place (grey line). The *second* panel shows the volumes at the individual bid and ask levels between prices of \$980.8 and \$982.8. Each unit on the x-axis is one message. The y-axis represents the price of Platinum in dollars. The colour represents the volume at each price level of the LOB for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the cumulative trade volume per second. The fourth panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the first spoof order of five contracts was placed, when the first contract of the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

second bid level, an upward staircase pattern emerged. After the first spoof order is cancelled, the same staircase pattern emerged but downwards. The height of the steps shows that the added and subtracted volumes were identical, that is, five contracts per step.

4.5.3 | Layered spoofing with iceberg orders: Visualization of cancellations around spoofing

Cancellations on the second bid and ask levels around the spoof in the Platinum July 2016 contract are visualized in Figure 18. During the visualized time window, zero contracts were cancelled on both the bid and ask side when the genuine order was placed. Between the first spoof order being placed and being cancelled, cumulative cancellations amounted to one contract on the bid side and three contracts on the ask side. Once the first spoof order was cancelled, another upward staircase pattern emerged on the bid side with identical heights of the steps, indicating that the cancellations had identical volumes. After all spoof orders from JPM were cancelled on the second bid level, cancellations continued in the visualized time




FIGURE 17 Visualization of second-level bid and ask volume behaviour around the spoof of 22 June 2016 in the Platinum July 2016 futures market. The *first* panel shows the volume of the second ask level. The *second* panel shows the volumes at the individual bid and ask levels between prices of \$980.8 and \$982.8. Each unit on the *x*-axis is one message. The *y*-axis represents the price of Platinum in dollars. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the volume of the second bid level. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the first spoof order of five contracts was placed, when the first contract of the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

window, albeit less frequently. The second level on the ask side showed no cancellations in the time window one second after the genuine order was executed.

4.5.4 | Layered spoofing with iceberg orders: Visualization of liquidity around spoofing

Figure 19 shows the ask and bid APM around the spoofing in the Platinum July 2016 market. Apart from one relatively large decrease, the liquidity costs on the bid side were relatively stable between 5 and 7 bps. Once the first spoof order was placed, liquidity costs decreased stepwise with each additional spoof order. Liquidity costs decreased from approximately 5.25 to 1.5 bps. Similarly, when the first spoof order was cancelled, liquidity costs increased stepwise with each spoof order cancelled. Ask side liquidity costs fluctuated between 8 and 10.5 bps during all JPM spoofing actions.

Table 12 shows the results of Welch's *t* tests used to test whether liquidity costs were significantly different before, during and after the spoof. For all different time windows, liquidity costs were higher before the spoof than during the spoof, meaning that liquidity increased during the spoof. Similarly, liquidity costs were lower during the spoof than after the spoof for all time windows. In other words,




FIGURE 18 Visualization of cumulative second-level bid and ask cancellation volume around the spoof of 22 June 2016 in the Platinum July 2016 futures market. The *first* panel shows the cumulative volume of cancellations of the second ask level. The *second* panel shows the volumes at the individual bid and ask levels between prices of \$980.8 and \$982.8. Each unit on the x-axis is one message. The y-axis represents the price of the Platinum in dollars. The colour represents the volume at each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the cumulative volume of cancellations of the second bid level. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the first spoof order of five contracts was placed, when the first contract of the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

liquidity was better during the spoof than after the spoof. Moreover, when comparing the liquidity costs before and after the spoof, liquidity was better before than after the spoof, as the liquidity costs after the spoof were higher than before.

4.6 | Liquidity as a motivation for spoofing

In previous sections, we proposed an alternative explanation for the use of spoofing, namely, attracting liquidity rather than moving the price. Table 13 summarizes for each spoofing example identified by the CFTC, whether our results correspond to the motivation of attracting more liquidity. The second column of Table 13, ‘Genuine order: placed on first level’, corresponds to the situation in which JPM seeks to attract more liquidity by placing the genuine order on the first bid or ask level—as, otherwise, they would have placed the genuine order deeper in the LOB and would have used the spoof to push the price through the first level(s) and hence get a better price than before. The third column of Table 13, ‘Genuine order: price identical to last traded price’, conforms to the situation

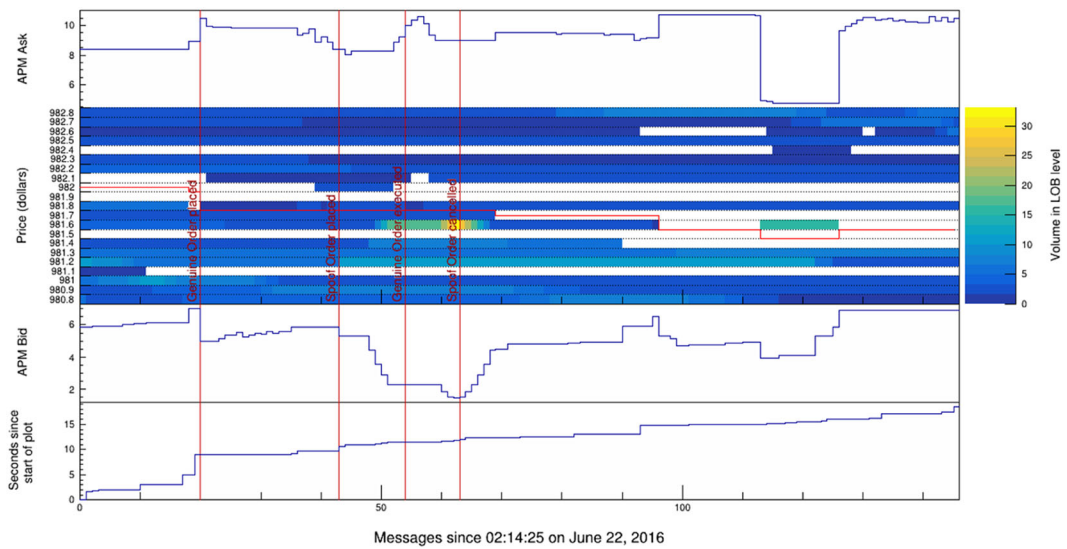


FIGURE 19 Visualization of bid and ask liquidity costs (Adverse Price Movement [APM]) behaviour around the spoof of 22 June 2016 in the Platinum July 2016 futures market. The *first* panel shows the APM of the ask side. APM measures the liquidity costs (in basis points) of a trader who wants to buy or sell a specific dollar value by submitting market orders. The *second* panel shows the volumes at the individual bid and ask levels between prices of \$980.8 and \$982.8. Each unit on the *x*-axis is one message. The *y*-axis represents the price of Platinum in dollars. The colour represents the volume in each price level of the limit order book (LOB) for each message. The scale ranges from blue to yellow, with the colour becoming a brighter yellow as volume increases at that price level. The red line is the midpoint. The *third* panel shows the APM for the bid side. The *fourth* panel shows how much time passes between messages reported by the exchange. A steeper (flatter) blue line signals a lower (higher) rate of messages, given that a steeper (flatter) line signals more (less) time progression. The red vertical lines signal when the JPMorgan spoofing activities took place, from left to right: when the genuine iceberg order was placed, when the first spoof order of five contracts was placed, when the first contract of the genuine order was executed and when the first spoof order was cancelled [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 12 Mean bid liquidity costs (bps) around the Platinum July 2016 spoof for different time windows

This table reports the mean liquidity costs (basis point [bps]), measured by Adverse Price Movement [APM] around the spoof in the Platinum July 2016 market for different periods and various time windows. *Before* represents the time up until the spoof order was added; *During* the period from when the spoof order was added until it was cancelled and *After* the time following the cancellation of the spoof order. Five different time windows are used, the *Spoof Duration* time window being 4.76 seconds. A lower APM indicates that liquidity costs are low and, hence, liquidity is high. Welch's *t* tests were used to test for mean differences between the periods. Significance at the 0.1% (two-tailed) level is indicated by ***.

Time window	Before versus during	During versus after	Before versus after
Spoof duration	5.969 > 4.419***	4.419 < 6.763***	5.969 < 6.763***
10 seconds	6.045 > 4.419***	4.419 < 6.829***	6.045 < 6.829***
30 seconds	6.402 > 4.419***	4.419 < 7.024***	6.402 < 7.024***
1 minute	6.948 > 4.419***	4.419 < 8.896***	6.948 < 8.896***
5 minutes	9.225 > 4.419***	4.419 < 10.543***	9.225 < 10.543***

TABLE 13 Liquidity as motivation for spoofing for each JPMorgan spoofing example

This table reports for each spoofing example three indicators for the motivation to use spoofing to attract liquidity. ‘Yes’ (‘No’) indicates that results conform (do not conform) to attracting liquidity. *Genuine order: placed on first level* indicates if the genuine order is placed on the first level. *Genuine order: price identical to last traded price* indicates if the price of the genuine order was identical to the last traded price. *Increase of liquidity after the spoof* shows if liquidity immediately after the spoof (*Spoof Duration*) was better than before the spoof.

Spoofing example	Genuine order: placed on first level	Genuine order: price identical to last traded price	Increase of liquidity after the spoof
Traditional spoofing			
10-Year T-Note December 2011	Yes	Yes	No
10-Year T-Note March 2010	Yes	No	No
Traditional spoofing with iceberg orders			
Ultra T-Bond September 2015	Yes	Yes	Yes
Silver March 2014	Yes	No	No
Layered spoofing			
Silver March 2012	Yes	No	No
Silver May 2014	Yes	Yes	Yes
Gold April 2014	Yes	No	Yes
T-Bond September 2009	No	No	Yes
Layered spoofing with iceberg orders			
Platinum July 2016	Yes	No	No

when the price of the genuine order is identical to the last traded price. The fourth column of Table 13, ‘Increase of liquidity after the spoof’, shows whether liquidity is better immediately after the spoof than before the spoof. We use the ‘Spoof Duration’ time window to determine this for each spoofing example.

Table 13 shows that there are cases in which attracting liquidity seems to be the motivation for spoofing. The Ultra T-Bond September 2015 and Silver May 2014 spoofing examples have all indicators point towards attracting liquidity as the motivation behind the spoof. In these cases, JPM was successful at attracting more liquidity: even after the spoof orders were cancelled, liquidity was higher after than before the spoof. Hence, JPM may have spoofed the market to keep prices stable and bait more traders into trading against their preferred price. In the other spoofing examples, one or two indicators confirm the motivation of attracting liquidity, that is, there is no spoofing example with all spoofing indicators being ‘No’.

5 | CONCLUSION

This study delved deeply into the JPM spoofing case and visualized their spoofing strategies from different angles. Using messages as its primary component, rather than time-based snapshots, a novel visualization methodology was used from particle physics to identify the JPM spoofing cases. This methodology allows researchers to study high-frequency data at a particular point in time (in our case, the time window of the spoofing), while also placing this data in the perspective of the market environment, that is, the entire LOB and related variables such as trades, bid and ask volumes, cancelled volume and liquidity. In other words, the message-based approach allows for the simultaneous visualization of activities in the LOB as well as surrounding activities, (re)actions and market output (e.g., price changes, liquidity). The time axis can be dynamically compressed or inflated to show the full details of the spoof, while leaving ample space for the state of the LOB before and after the spoofing activities. This visualization method (1) shows how well-hidden spoofing can be; (2) provides insights in the complexity of the techniques required to recognize spoofing and (3) puts a value on the minuscule price changes that make spoofing economically viable. We analyse the JPM spoofing examples as identified by the CFTC in detail with numerous characteristics and, in some cases, propose an alternative explanation of why JPM spoofed the market. Rather than move the market to their benefit by inducing short-term price trends, their intention may sometimes have been to attract liquidity, so as to buy or sell numerous futures contracts without having to bear the financial consequences of an illiquid market (i.e., incur costs for trading in a less-than-perfectly liquid market). These visualizations offer a glimpse of the patterns, techniques, time scales, and motivations of the spoofer, thus yielding invaluable information for fraud detection. Messages are visualized in a unique way and help to retrieve more retrospective information about patterns in the LOB at the time when a trader spoofed the market. Reconstructing and visualizing the LOB is key to detecting spoofing, as raw data presents an incomplete overview that does not show orders or changes in the market in relation to its context. Environmental and contextual variables are needed to understand order and market behaviour as a whole. However, the data and visualizations alone are not sufficient to identify (new types of) spoofing.

Gained spoofing insights and visualizations have implications for all stakeholders. Both academics and industry participants gain a better understanding of various types of spoofing and how the market behaves during spoofing. New insights into the motives of market manipulation will help academics to model market behaviour in, for example, agent-based modelling. The provided visualization demonstrates how high-frequency LOB data can be effectively visualized and why message-based visualizations contain more information than time-based visualizations. Both academics and industry participants can use these visualizations and adjust them to any variable of interest. Regulators and exchanges gain a different perspective on spoofing as they can now observe all market activity, rather than have to resort to aggregated market activity. Moreover, the visualizations can enhance and refine surveillance programs.

The visualization approach in this study may encourage and inspire future researchers to use more diverse LOB visualization methodologies. Future research might focus on which types of spoofing can be visualized and which go undetected. Moreover, large portions of trading in equity markets are nowadays driven by algorithms. Future research could examine how visualizations may help to control potential spoofing activities by algorithmic trading. Also, the proposed visualization allows for an alternative explanation of spoofing as a means

to attract liquidity. We did not know the true intentions of JPM and can only speculate on their intentions. To further examine the motivation of spoofing, further research can focus on in-depth interviews; behavioural and experimental studies to identify the set of motivations for spoofing and the relationship between spoofing and liquidity costs. Furthermore, this study may motivate future research into the development of theoretical frameworks that can help us to better understand anomalies and market manipulation in financial markets. Finally, the use of iceberg orders in spoofing may trigger a debate about the visibility of orders to regulators and market participants. Future research may have to address whether the use of iceberg orders is fair, whether these orders facilitate manipulative practices, such as spoofing, and whether it still makes sense to allow them in a modern trading environment with algorithmic traders. The message-based visualizations proposed in this study may contribute to this debate.

ACKNOWLEDGMENTS

The authors would like to thank the European Organization for Nuclear Research (CERN, Geneva) for extending their analytical ROOT system and providing data storage and computing power for this research. We would like to thank the Commodity Risk Management Expertise Centre (CORMEC) for providing their infrastructure. We are grateful to the Chicago Mercantile Exchange Foundation for providing market-depth data for all JPMorgan spoofing examples as identified by the CFTC, containing all market messages, timestamps and levels of the limit order book. We would like to thank the Office for Futures and Options Research at the University of Illinois at Urbana-Champaign, the Dutch Authority for the Financial Markets (AFM), Euronext and the Dutch National Bank (DNB) for their constructive feedback and insights on previous versions of this manuscript. We would like to thank SURF SARA for granting access to the national computing and data-storage environment (Grud HPC Cloud Beehub, Hadoop).

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the Chicago Mercantile Exchange. Restrictions apply to the availability of these data, which were used under license for this study. The data are available from the authors with the permission of the Chicago Mercantile Exchange.

ORCID

Philippe Debie  <https://orcid.org/0000-0002-7353-9715>

Cornelis Gardebroek  <https://orcid.org/0000-0002-3154-9464>

Stephan Hageboeck  <https://orcid.org/0000-0001-9359-2196>

Axel Naumann  <https://orcid.org/0000-0002-4725-0766>

Andres A. Trujillo-Barrera  <https://orcid.org/0000-0002-2740-8327>

Marjolein E. Verhulst  <https://orcid.org/0000-0003-4134-2288>

REFERENCES

- Aidov, A., & Daigler, R. T. (2015). Depth characteristics for the electronic futures limit order book. *The Journal of Futures Markets*, 35(6), 542–560. <https://doi.org/10.1002/fut.21706>
- Antcheva, I., Ballintijn, M., Bellenot, B., Biskup, M., Brun, R., Buncic, N., Canal, P., Casadei, D., Couet, O., Fine, V., Franco, L., Ganis, G., Gheata, A., Maline, D. G., Goto, M., Iwaskiewicz, J., Kreshuk, A.,

- Segura, D. M., Maunder, R., ... Tadel, M. (2009). ROOT—A C++ framework for petabyte data storage, statistical analysis and visualization. *Computer Physics Communications*, 180(12), 2499–2512. <https://doi.org/10.1016/j.cpc.2009.08.005>
- Battalio, R., Corwin, S. A., & Jennings, R. (2016). Can brokers have it all? On the relation between make-take fees and limit order execution quality. *The Journal of Finance*, 71(5), 2193–2238. <https://doi.org/10.1111/jofi.12422>
- Biais, B., Bisière, C., & Spatt, C. (2010). Imperfect competition in financial markets: An empirical study of Island and Nasdaq. *Management Science*, 56(12), 2237–2250. <https://doi.org/10.1287/mnsc.1100.1243>
- Bongiovanni, S., Borkovec, M., & Sinclair, R. D. (2006). Let's play hide-and-seek. *The Journal of Trading*, 1(3), 34–46. <https://doi.org/10.3905/jot.2006.644087>
- Brogaard, J. (2010). *High frequency trading and its impact on market quality* (Working paper no. 66). Northwestern University Kellogg School of Management.
- Brogaard, J., & Garriott, C. (2019). High-frequency trading competition. *Journal of Financial and Quantitative Analysis*, 54(4), 1469–1497. <https://doi.org/10.1017/S0022109018001175>
- Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. *Review of Financial Studies*, 27(8), 2267–2306. <https://doi.org/10.1093/rfs/hhu032>
- Brogaard, J., Hendershott, T., & Riordan, R. (2019). Price discovery without trading: Evidence from limit orders. *The Journal of Finance*, 74(4), 1621–1658. <https://doi.org/10.1111/jofi.12769>
- Brun, R., & Rademakers, F. (1997). ROOT—An object oriented data analysis framework. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389(1–2), 81–86.
- Buti, S., & Rindi, B. (2013). Undisclosed orders and optimal submission strategies in a limit order market. *Journal of Financial Economics*, 109(3), 797–812. <https://doi.org/10.1016/j.jfineco.2013.04.002>
- Cartea, Á., Jaimungal, S., & Wang, Y. (2020). Spoofing and price manipulation in order-driven markets. *Applied Mathematical Finance*, 27(1–2), 67–98. <https://doi.org/10.1080/1350486X.2020.1726783>
- Chakrabarty, B., & Shaw, K. W. (2008). Hidden liquidity: Order exposure strategies around earnings announcements. *Journal of Business Finance & Accounting*, 35(9–10), 1220–1244. <https://doi.org/10.1111/j.1468-5957.2008.02111.x>
- Chordia, T., Hu, J., Subrahmanyam, A., & Tong, Q. (2019). Order flow volatility and equity costs of capital. *Management Science*, 65(4), 1520–1551. <https://doi.org/10.1287/mnsc.2017.2848>
- Colliard, J.-E., & Hoffmann, P. (2017). Financial transaction taxes, market composition, and liquidity. *The Journal of Finance*, 72(6), 2685–2716. <https://doi.org/10.1111/jofi.12510>
- Commodity Futures Trading Commission (CFTC). (2020). *Order instituting proceedings pursuant to Section 6(c) and (d) of the commodity exchange act, making findings, and imposing remedial sanctions*, Washington, D.C. Release Number: 8260-20. <https://www.cftc.gov/media/4826/enfjpmorganchaseorder092920/download>
- Cont, R., Kukanov, A., & Stoikov, S. (2014). The price impact of order book events. *Journal of Financial Econometrics*, 12(1), 47–88. <https://doi.org/10.1093/jfinec/nbt003>
- Dacorogna, M. M., Gençay, R., Müller, U. A., Olsen, R. B., & Pictet, O. V. (2001). an introduction to high-frequency finance. <https://doi.org/10.1016/B978-0-12-279671-5.X5000-X>
- Dalko, V., & Wang, M. H. (2018). How effective are the order-to-trade ratio and resting time regulations? *Journal of Financial Regulation*, 4(2), 321–325. <https://doi.org/10.1093/jfr/fjy007>
- Dalko, V., & Wang, M. H. (2020). Volume limit: An effective response to the India flash crash? *Journal of Financial Regulation*, 5(2), 249–255. <https://doi.org/10.1093/jfr/fjz006>
- Dugast, J. (2018). Unscheduled news and market dynamics. *The Journal of Finance*, 73(6), 2537–2586. <https://doi.org/10.1111/jofi.12717>
- European Organization for Nuclear Research (CERN). (2018a). About ROOT. Retrieved November 25, 2020, from <https://root.cern.ch/about-root>
- European Organization for Nuclear Research (CERN). (2018b). ROOT. <https://root.cern>
- Fabozzi, F. J., Focardi, S. M., & Jonas, C. (2011). High-frequency trading: Methodologies and market impact. *Review of Futures Markets*, 9(Special Issue), 7–38.
- Gomber, P., & Schweickert, U. (2002, July). *The market impact—Liquidity measure in electronic securities trading*. Die Bank. <https://www.xetra.com/resource/blob/5480/784401adc4e0e19bf32e37edcb45926d/data/The-Market-Impact-Liquidity-Measure-in-Electronic-Securities-Trading.pdf>

- Gomber, P., Schweickert, U., & Theissen, E. (2015). Liquidity dynamics in an electronic open limit order book: An event study approach. *European Financial Management*, 21(1), 52–78. <https://doi.org/10.1111/j.1468-036X.2013.12006.x>
- Hasbrouck, J. (2018). High-frequency quoting: Short-term volatility in bids and offers. *Journal of Financial and Quantitative Analysis*, 53(2), 613–641. <https://doi.org/10.1017/S0022109017001053>
- Hautsch, N., & Horvath, A. (2019). How effective are trading pauses? *Journal of Financial Economics*, 131(2), 378–403. <https://doi.org/10.1016/j.jfineco.2017.12.011>
- Ito, T., & Yamada, M. (2018). Did the reform fix the London fix problem? *Journal of International Money and Finance*, 80, 75–95. <https://doi.org/10.1016/j.jimonfin.2017.10.004>
- Kahraman, B., & Tookes, H. E. (2017). Trader leverage and liquidity. *The Journal of Finance*, 72(4), 1567–1610. <https://doi.org/10.1111/jofi.12507>
- Leangarun, T., Tangamchit, P., & Thajchayapong, S. (2016). Stock price manipulation detection based on mathematical models. *International Journal of Trade Economics and Finance*, 7(3), 81–88. <https://doi.org/10.1109/ICACI.2016.7449848>
- Lee, E. J., Eom, K. S., & Park, K. S. (2013). Microstructure-based manipulation: Strategic behavior and performance of spoofing traders. *Journal of Financial Markets*, 16(2), 227–252. <https://doi.org/10.1016/j.finmar.2012.05.004>
- Linton, O., & Mahmoodzadeh, S. (2018). Implications of high-frequency trading for security markets. *Annual Review of Economics*, 10(1), 237–259. <https://doi.org/10.1146/annurev-economics-063016-104407>
- Mendonça, L., & de Genaro, A. (2020). Detection and analysis of occurrences of spoofing in the Brazilian capital market. *Journal of Financial Regulation and Compliance*, 28(3), 369–408. <https://doi.org/10.1108/JFRC-07-2019-0092>
- Menkveld, A. J., & Yueshen, B. Z. (2019). The flash crash: A cautionary tale about highly fragmented markets. *Management Science*, 65(10), 4470–4488. <https://doi.org/10.1287/mnsc.2018.3040>
- Michaels, D. (2020, September 29). JPMorgan paying \$920 million to resolve market manipulation probes. *Wall Street Journal*. [ps://www.wsj.com/articles/jpmorgan-paying-920-million-to-resolve-market-manipulation-probes-11601393666](https://www.wsj.com/articles/jpmorgan-paying-920-million-to-resolve-market-manipulation-probes-11601393666)
- Neurensic. (2016). Spoofing similarity model. Retrieved October 13, 2020, from <https://neurensic.com/spoofing-similarity-model/>
- Paddrik, M. E., Haynes, R., Todd, A. E., Scherer, W. T., & Beling, P. A. (2016). Visual analysis to support regulators in electronic order book markets. *Environment Systems and Decisions*, 36(2), 167–182. <https://doi.org/10.1007/s10669-016-9597-2>
- Putniņš, T. J. (2012). Market manipulation: A survey. *Journal of Economic Surveys*, 26(5), 952–967. <https://doi.org/10.1111/j.1467-6419.2011.00692.x>
- Schoenberg, T., & Robinson, M. (2020, September 29). JPMorgan admits spoofing by 15 traders, two desks in record deal. *Bloomberg*. <https://www.bloomberg.com/news/articles/2020-09-29/jpmorgan-pays-920-million-admits-misconduct-in-spoofing-probe>
- Sensoy, A. (2019). Commonality in ask-side vs. bid-side liquidity. *Finance Research Letters*, 28, 198–207. <https://doi.org/10.1016/j.frl.2018.04.020>
- Tejedor, E., & Kothuri, P. (2018). CERN's platform for data analysis with spark. Retrieved October 18, 2020, from <https://www.slideshare.net/databricks/cerns-next-generation-data-analysis-platform-with-apache-spark-with-enric-tejedor>
- United States. (2010). Dodd-Frank Wall Street Reform and Consumer Protection Act: Conference Report (to Accompany HR 4173).
- Verhulst, M. E., Debie, P., Hageboeck, S., Pennings, J. M. E., Gardebroek, C., Naumann, A., Leeuwen, P., Trujillo-Barrera, A. A., & Moneta, L. (2021). When two worlds collide: Using particle physics tools to visualize the limit order book. *Journal of Futures Markets*, 41(11), 1715–1734.
- Yao, C., & Ye, M. (2018). Why trading speed matters: A tale of queue rationing under price controls. *The Review of Financial Studies*, 31(6), 2157–2183. <https://doi.org/10.1093/rfs/hhy002>
- Zhai, J., Cao, Y., Yao, Y., Ding, X., & Li, Y. (2017). Computational intelligent hybrid model for detecting disruptive trading activity. *Decision Support Systems*, 93, 26–41. <https://doi.org/10.1016/j.dss.2016.09.003>

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Debie, P., Gardebroek, C., Hageboeck, S., Leeuwen, P. v., Moneta, L., Naumann, A., Pennings, J. M. E., Trujillo-Barrera, A. A., & Verhulst, M. E. (2023). Unravelling the JPMorgan spoofing case using particle physics visualization methods. *European Financial Management*, 29, 288–326.

<https://doi.org/10.1111/eufm.12353>