HAL open science

Multiple-limit trades : empirical facts and application to lead-lag measures

Fabrizio Pomponio, Frédéric Abergel

To cite this version:

Fabrizio Pomponio, Frédéric Abergel. Multiple-limit trades : empirical facts and application to lead-lag measures. Quantitative Finance, 2013, 13 (5), pp.783-793. 10.1080/14697688.2012.743671 . hal-00745317

HAL Id: hal-00745317

https://hal.science/hal-00745317v1

Submitted on 25 Oct 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Quantitative Finance, Vol. 00, No. 00, Month 200x, 1–23

Multiple-limits trades : empirical facts and application to lead-lag measures.

FABRIZIO POMPONIO*†‡ and FREDERIC ABERGEL†

†BNP Paribas Chair of Quantitative Finance, MAS Laboratory, Ecole Centrale Paris, Grande Voie des Vignes, Chatenay-Malabry, 92290, France
‡ BNP Paribas, Equities & Derivatives Quantitative R&D, 20 boulevard des Italiens, Paris, 75002, France
(Received 00 Month 200x; in final form 00 Month 200x)

Order splitting is a standard practice in trading : traders constantly scan the limit order book and choose to limit the size of their market orders to the quantity available at the best limit, thereby controlling the market impact of their orders. In this article, we focus on the other trades, multiple-limits trades that go through the best available price in the order book, or "trade-throughs". We provide various statistics on trade-throughs: frequency, volume, intraday distribution, market impact... and present a new method for the measurement of lead-lag parameters between assets, sectors or markets.

Keywords: Financial markets; Market microstructure; High-frequency trading

1. Introduction

It is a well-documented fact that the order flow is a highly autocorrelated long-memory process, see e.g. Bouchaud *et al.* (2004), Lillo and Farmer (2004) [1]. In Lillo and Farmer (2004), the empirical autocorrelation function of trades signs for 20 stocks traded at the London Stock Exchange is fitted to a power-law $C(\tau) \sim \tau^{-\gamma}$ with exponent $\gamma = 0.6$. On the Paris Stock Exchange, Bouchaud *et al.* (2004) measure power-law exponents γ ranging from 0.2 to 0.7. In other words, the sign of a trade has an impact on futures trades signs, with levels that stay statistically significant over a period of time as long as two weeks.

Different models have been suggested in the literature to explain this phenomenon, see the review in Bouchaud *et al.* (2009). LeBaron and Yamamoto (2007) consider an evolutive market of heterogeneous investors allowed to learn and adapt their trading, and manage to replicate the long memory of order signs thanks to the imitative behavior of investors. Lillo *et al.* (2005) argue that long memory is mainly due to a delay in market clearing. A trader facing a large

*Corresponding author. Email: fabrizio.pomponio@gmail.com

[1]A centered process X is said to exhibit long-memory behavior when its series of auto-covariances is not summable, i.e. $\sum_{h\in\mathbb{N}} |\gamma(h)| = +\infty$ where $\gamma(h) = E[X_t X_{t-h}]$

Quantitative Finance
ISSN 1469-7688 print/ISSN 1469-7696 online © 200x Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/1469768YYxxxxxxxx

order will split it in several orders for two main reasons: the available liquidity in the order book may not be sufficient and, even if it were, revealing her/his intention to the market will cause the price to move too much in the wrong direction.

Empirical statistics that discriminate between those two explanations are not easy to make, because of the lack of available data on individual investors. However, some exchanges (like the London Stock Exchange, the Spanish Stock Exchange and the Australian Stock Exchange) provide a partial information (the membership code) that identify the member of the exchange who executes the trade. Gerig (2007) computes the empirical autocorrelation function of orders signs considering market orders coming from the same membership code, different membership codes, and all market orders. When conditioning on the same membership code, the autocorrelation function is one order of magnitude bigger than the unconditional one, and is also decaying as a power-law, but with a smaller exponent. On the contrary, when conditioning on different membership codes, the empirical autocorrelation function of orders signs is fastly decaying to zero and is not a power-law anymore. If we assume that most investors trade only through a small number of brokers, these statistics support the explanation of Lillo *et al.* (2005) rather than that of LeBaron and Yamamoto (2007): the long-memory of orders signs is due to delay in market clearing.

So, in practice, the splitting of orders is an important pattern of market microstructure: traders wanting to trade a large order constantly scan the limit order book and split their orders to restrict their size to the quantity available at the best limit. However, speed is sometimes more important than minimizing market-impact. In this article, we study trade-throughs, which are precisely the trades that stand outside the usual trading pattern.

The statistical properties of limit order books have been extensively studied, see for example Bouchaud and Potters (2003), Bouchaud *et al.* (2009) and the references therein. We will revisit some well-known statistics of limit order books, albeit restricted to trade-throughs. For example, we will study whether market impact (the average response of prices to trades), which has received considerable attention over the last years, see e.g. Potters and Bouchaud (2003), Weber and Rosenow (2005), Almgren *et al.* (2005), Hautsch and Huang (2009), Eisler *et al.* (2009), Cont *et al.* (2011), changes when one considers the response to trade-throughs.

For obvious reasons that become even clearer in the bulk of the paper, trade-throughs are naturally related to large trades and large price moves. Previous research works such as Farmer *et al.* (2004) or Weber and Rosenow (2006) investigate large moves in limit order books, mainly focusing on their cause: are they due to large trade volume, or lack of liquidity ? Farmer *et al.* (2004) found that large price changes on the London Stock Exchange occur when a market order removes all the volume at the best limit (as trade-throughs do), thereby creating a change in the corresponding best price equal to the size of the first gap. This indicates that a major contribution of large price changes is due to fluctuations in the liquidity inside the limit order books. Similar results have been obtained by Weber and Rosenow (2006) on NASDAQ data and confirmed that extreme price fluctuations were mainly caused by a small liquidity in the limit order book. In this paper, we somehow make this interpretation more complete, by studying the rather peaked distribution of the occurrences of trade-throughs, connecting them with exogenous events such as market opening and closing or economic news announcements.

Also of great practical interest is the study of lead-lag relationships between two assets or market places. A lead-lag relationship occurs when knowledge of the return of one asset helps predict the future return of a second asset. Lead-lag relationships have received much attention in the recent literature, see Abergel and Huth (2011) and Hoffmann and Rosenbaum (2011) and the references therein. A standard approach is to measure the lagged cross-correlation of the returns, and to study a possible asymetry between positive and negative lags, or even, to find a non-zero lag that maximizes the cross-correlation. Our approach is different, and we believe it

is new. We try and answer the following question: which of the two assets moves first ? In other words, when an important move occurs on Asset 1, is it followed or is it caused by an important move on Asset 2 ? When focusing only on trade-throughs as we do, and due to the relative rarity of such events, we will show that one can answer that question in a statistically significant way.

The paper is organized as follows: in Section 2, we give a general presentation of the TRTH database used in this article. We also present the set of data we focus on (major US and EU equity futures and major French stocks). In Section 3, we give a precise definition of a trade-through and study some elementary statistical properties [1]: arrival time distribution, volume, seasonality, clustering. Section 4 is devoted to the study of the market impact of trade-throughs. In Section 5, we examine the behaviour of the spread after a trade-through, showing a typical power-law relaxation of the excess spread. Finally, Section 6 deals with the characterization of lead-lag relationships between US and EU equity markets and between pairs of French stocks of the same sector.

[1] All computations are made using the free statistical software `R`, available at http://cran.r-project.org.

Table 1. TRTH's most important trades flags

Flag	Information and signification
normal	Trades occurring during the continuous trading session.
auction	Trades occurring at the end of the auction phase.
OTC	Over-the-counter trades occurring directly between two parties. They are opposed to trades occurring on centralized exchanges.
off book	Trades occurring outside the usual trading system of the considered exchange. They may be trades reported from a broker, for example. In Euronext, trades outside of the NSC (Nouveau Système de Quotation) fall inside this category.
block trade	Reporting of block trades (very large trades).
rck	TRTH/exchange's threshold break alerting that price and/or volume of this trade (as reported by the exchange) seem too different from an usual behavior and should not be considered as relevant.
market closed	Trades occuring before or after the regular trading session.
cancelled	Cancelled trades.
late0day & lateNdays	Trades that were reported the same day or N days after.
late report	Trades reported later in market data feeds. They include late0day and lateNdays trades.
unknown	All the other trades.

2. Data presentation

2.1. General TRTH data presentation and processing

The data used in this study come from the TRTH (Thomson Reuters Tick History) databases. There are two different databases, one for the quotes (grouped in the 'Quotes' file) and one for the trades (grouped in the 'Time And Sales' file). Both quotes and trades are timestamped in milliseconds by Thomson-Reuters timestamping machines located in London. Quotes entries are composed of Bid/Ask/BidSize/AskSize. Trades entries contain Price/Volume of each transaction.

An important point to be mentioned in the data presentation is that TRTH data are flagged. Each entry of both quotes and trades files has a flag indicating information to be taken into account in the data analysis. Those flags are market- and exchange-dependent in the sense that specific knowledge of each market and exchange is necessary to correctly interpret each TRTH/exchange flag. After processing the flags [1], we end up with trades tagged within a limited number of flags' categories. The most important trades flags (normal, auction, OTC, offbook, block trade, rck, market closed, cancelled, late0day & lateNdays, late report) are detailed in Table 1.

Note that some data sent by exchanges are corrections of previous entries, such as the cancellation of a previous trade and its replacement by another one. Corrections are the only case when TRTH data are modified at BNP Paribas before being made accessible to users.

2.2. Data used in this study

Tick-by-tick market data used in this study are the data after all corrections have been taken into account. We consider only trades that are flagged as 'normal' trades. In particular, we do

[1]The flag processing is done by BNP Paribas Equities & Derivatives Quantitative R&D Histo team.

not consider any block-trade or off-book trade. Moreover, we restrict ourselves to data coming from a single venue, the main exchange where the considered assets are traded.

2.2.1. EU-US equity futures

Here is how we choose the data for the comparison of the US and EU equity market: first, we select the most representative instruments in the US and EU equity markets. To this end, we rank all equity financial instruments available in TRTH according to their ADV (Average Daily Volume), and then pick the most liquid ones (3 from US equity markets and 3 from EU equity markets). By doing so, we end up with a small number of financial instruments representing the most liquid instruments of the markets under scrutiny. This choice is based on the rationale that market moves are first seen in the most liquid instruments, as those most liquid assets tend to incorporate information faster.

This selection process provides a set composed of E-mini S&P500, Nasdaq E-mini and Dow Jones E-mini futures (for US equity markets) and Eurostoxx, DAX and Footsie futures (for European equity markets). These assets trade on the CME Exchange (for E-mini S&P500, Nasdaq E-mini and Dow Jones E-mini futures), the Eurex Exchange (for Eurostoxx and DAX futures) and the NYSE Liffe Exchange (for the Footsie future). Moreover, in order to consider the most liquid instruments, we focus on the futures with the nearest maturity.

For the basic statistics on frequency and volumes of trade-throughs, we use the data of March 2010 (from 16 to 21.30, Paris Time reference). For the intraday distribution, we focus on the first half of March 2010 (to avoid difference in daylight saving time between Paris and the USA). For the lead-lag study, we use data from the beginning of December 2009 to mid-March 2010 and restrict our data time-frame to the period of the day when both EU and US equity markets are open and widely trading (from 15.30 to 17.30, Paris Time reference), as the lead-lag phenomenon is particularly relevant at that moment.

2.2.2. French stocks

For French stocks, the assets we select are BNP Paribas, Société Générale, Renault and Peugeot during March 2010, with a daily time frame from 9.30 to 17.00, Paris Time reference, so as not to be impacted by auction phases.

3. Elementary statistical properties of trade-throughs

Traders usually scan the limit order book and restrict the size of their orders to the available liquidity. If necessary, they split a large order into several smaller orders to control the trade size. As explained in the introduction, we are interested in the trades that deviate from this usual behavior. We focus on trades that consume the liquidity available in the order book in an aggressive way, namely the *trade-throughs*.

We define as an x-th limit trade-through, any trade that consumes at least one share at the x-th limit in the order book. For example, a 2nd-limit trade-through completely consumes the first limit and begins to consume the second limit of the order book. In Figure 1, we show an example of such a trade.

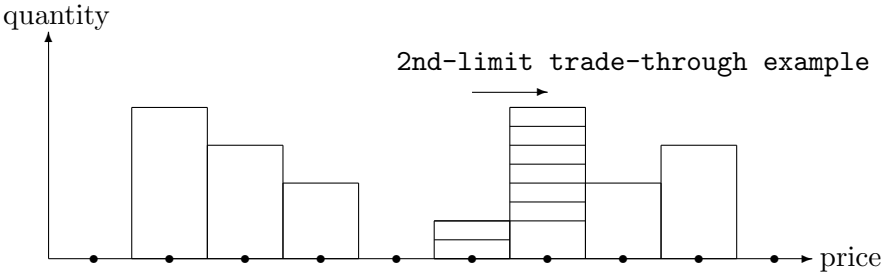

Figure 1. 2nd-limit trade-through example

Table 2. Basic statistics on 2nd-limit trade-throughs (estimations based on March 2010 data)

Financial asset considered	2nd-limit TT Occurrence (in %— daily number)	2nd-limit TT Volume (in %)	Relative tick value (indicative, in bp)
E-mini S&P500 - ES@	1,43—638	2,66	2,2
Nasdaq E-mini - NQ@	1,91—404	5,73	1,3
Dow Jones E-mini - YM@	3,17—479	9,65	0,9
Eurostoxx - STXE@	2,2—237	5,4	3,5
Footsee - FFI@	4,3—419	12,7	0,9
DAX - FDX@	8,57—724	22,6	0,8
Peugeot - PEUP.PA	4,3 — 112	12,2	2,3
BNP Paribas -BNPP.PA	4,3 — 298	11,3	1,8
Renault - RENA.PA	5,0 — 224	12,8	1,5
Société Générale - SOGN.PA	6,0 — 451	14,8	1,1

3.1. On which limit of the order book is liquidity taken from ?

We study the location where liquidity is taken from the order book. More precisely, we want to measure the fraction of the total trade volume that is taken from each limit of the order book. In Figure 2, we plot this fraction against the algebraic limit number (strictly positive for the ask side, strictly negative for the bid side and zero for the trades we could not find on which limit they occurred) for some stocks. For example, approximately 7% of the trading volume is taken from the first limit of the order book for the BNP Paribas stock (84% of which is taken from the second limit). This is clearly non-negligible and confirms that traders may sometimes consume the liquidity in a more aggressive way, rather than wait for new liquidity to be provided.

3.2. Frequency and volume

Here, we present basic statistics on occurrences and volumes of trade-throughs, in order to better measure the significance of this phenomenon.

Note that even if trade-throughs are rare events, they form a sizeable part of the daily volume (up to 20% for the DAX index future).

An important remark should be made at this stage: the smaller the relative tick value (which is the absolute tick value divided by the value of the asset), the more important trade-throughs are, both in occurrence and volume. This result seems natural in the sense that, the smaller the relative tick value on an asset is, the more aggressively this asset is traded.

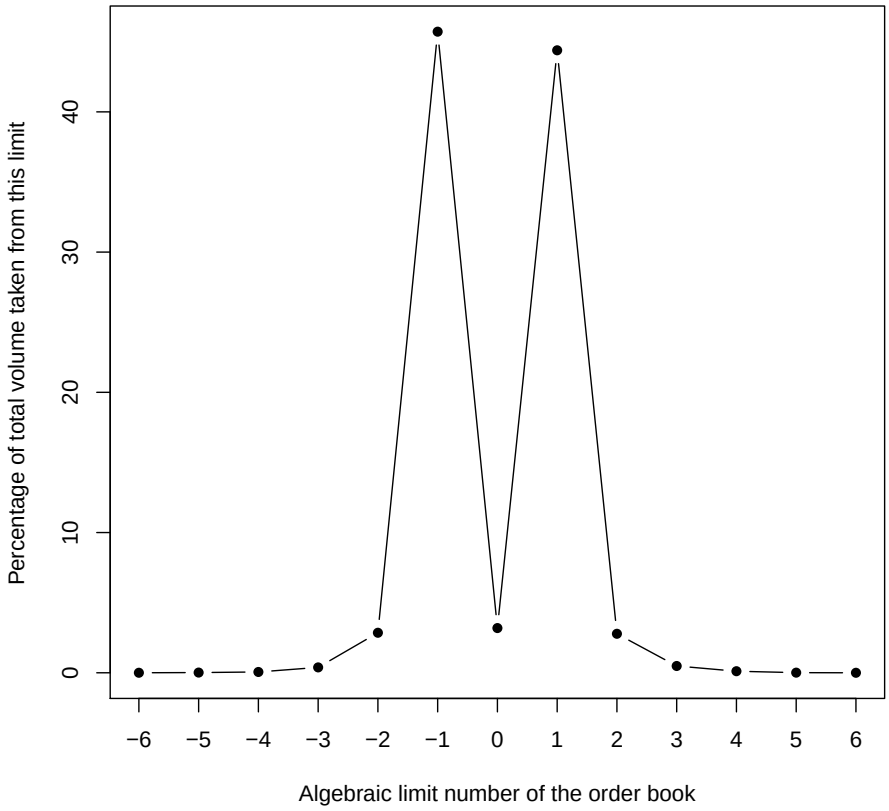

Figure 2. Fraction of total trading volume taken from each limit of the order book of BNP Paribas stock

Table 3. Basic statistics on 3rd-limit trade-throughs (estimations based on March 2010 data)

Financial asset considered	3rd-limit TT Occurrence (in bp— daily number)	3rd-limit TT Volume (in %)	Relative tick value (indicative, in bp)
E-mini S&P500 - ES@	0,12 — 0,6	0,0067	2,2
Nasdaq E-mini - NQ@	1,4—2,9	0,26	1,3
Dow Jones E-mini - YM@	6,2—8,5	0,79	0,9
Eurostoxx - STXE@	0,41 — 0,47	0,096	3,5
Footsie - FFI@	27—26	2,3	0,9
DAX - FDX@	49—42	3,58	0,8
Peugeot - PEUP.PA	59 — 16	3,47	2,3
BNP Paribas -BNPP.PA	48 — 33	2,49	1,8
Renault - RENA.PA	81 — 37	3,8	1,5
Société Générale - SOGN.PA	127—94	5,2	1,1

From a liquidity point of view, we also know that limit order books of assets with a big relative tick value tend to be more filled with high liquidity at the best, whereas order books of assets with a small relative tick value are less filled with liquidity at the best and may also present gaps.

Table 4. Stability of the definition of trade-throughs with respect to the Mean Trade Volume (estimations based on March 2010 data)

Financial asset considered	2nd-limit TT Occurrence fraction (in %)	3rd-limit TT Occurrence fraction (in %)	Relative tick value (indicative, in bp)
E-mini S&P500 - ES@	31	84	2,2
Nasdaq E-mini - NQ@	62	91	1,3
Dow Jones E-mini - YM@	69	90	0,9
Eurostoxx - STXE@	40	94	3,5
Footsee - FFI@	70	97	0,9
DAX - FDX@	80	99	0,8
Peugeot - PEUP.PA	75	96	2,3
BNP Paribas -BNPP.PA	74	96	1,8
Renault - RENA.PA	78	96	1,5
Société Générale - SOGN.PA	76	95	1,1

This explains why assets with small relative tick values are more often hit by trade-throughs. So, this liquidity pattern for assets with small relative tick values is consistent with the larger number of occurrences of trade-throughs in Tables 2 and 3 on those assets.

The typical daily number of such 2nd-limit trade-throughs is about a few hundreds (from about 100 for a French stock like Peugeot to about 600 for the E-mini S&P500 future). When looking at the statistics on 3rd-limit trade-throughs, we can see that this daily number falls to less than 10 events for US futures (which shows that almost all trade-throughs on those assets only reach the 2nd-limit of the order book), to less than 50 events for EU futures and 100 events for French stocks.

On the volume part, we can see that except for US futures and the Eurostoxx future, the volume fraction due to 3rd-limit trade-throughs is in the range of 2 % to 5%. Clearly, such an analysis shows that trade-throughs, and especially 2nd-limit trade-throughs, are significant as a fraction of the number of trades and the trade volume.

Let us now check that our definition of trade-throughs is stable with respect to the volume of the trade. To do so, we first compute the mean trade volume of the usual trades, i.e., trades that consume less than the quantity available at the first limit. Then, we check if the statistical set of data we defined as trade-throughs is stable when adding a volume condition to their definition, namely, that they consume more than the mean trade volume.

In the case of French stocks in Table 4, a fraction of 75% of the data considered as 2nd-limit trade-throughs is stable with the new definition, and this fraction goes to 95% for 3rd-limit trade-throughs.

For the US and EU futures, the situation is more complex. If we look only at the 3rd-limit trade-throughs, the fraction of trade-throughs stable with the new definition is over 84%. But for 2nd-limit trade-throughs, there exists a bigger difference in this fraction, especially for the E-mini S&P500 and the Eurostoxx futures (which are both under 40%).

For these two particular assets, a fraction of 60 to 70% of 2nd-limit trade-throughs are caused by relatively small trades (with a volume smaller than the mean trade volume of the usual trades). We notice that they have the biggest relative tick values among the considered futures. As the best limit is crossed by those trades, this means that trade-throughs on futures with big relative ticks tend to happen in regimes when liquidity is lacking on the best limits.

In conclusion, except for these two particular futures in the case of 2nd-limit trade-throughs, the definition is stable when we add a volume condition.

Table 5. Rarefaction of data for French stock in the estimation of conditional probability of trade-throughs

Volume of the trade (in MTV)	Proportion of trades with a volume higher than (in %)			
	BNPP.PA	SOGN.PA	RENA.PA	PEUP.PA
1	29	30	32	31
2	12	11	12	12
5	2	2	2	2
10	0.4	0.3	0.4	0.4
20	0.07	0.06	0.06	0.05

3.3. Links between trade-throughs and large trades

In this section, we examine the link between trade-throughs and large trades. Intuitively, the larger a trade, the more probable the fact that this trade is a trade-through. On Figure 3, is shown the conditional probability for a trade to be a trade-through if its volume is higher than a given threshold, where this threshold is expressed in multiples of the mean trade volume. For French stocks, this conditional probability roughly increases from 0 to 0.4 for trade volumes ranging from 0 to 20 mean trade volumes.

This shows that trade-throughs are related to big trades as expected. However, empirically, this relation is not as clear as one might expect, since conditional probabilities are not reaching levels close to 1 which would reflect that every big trade is a trade-through. A statistical flaw may explain this difference : the bigger the trade volume limit used to measure the conditional probability, the less data we have (as shown in Table 5). Hence, the most reliable part of the graph is that close to the origin, with a rapidly increasing conditional probability for a trade to be a trade-through as its volume increases. But even taking this issue into account, we believe based on our data that some differences between trade-throughs and large trades will remain, as that some very large trades simply reflects the presence of large quantities at the best limit.

3.4. Clustering

Clustering of trade-throughs is studied by looking at the arrival time of the next trade-through: if there is a clustering of trade-throughs, the next trade-through should arrive faster after a trade-through than after any trade. To verify this, we compute the empirical arrival time distribution of the next trade-through, conditioned or not by the fact that the current trade is a trade-through.

We can see on Figure 4 that the distribution is more peaked for short waiting-times (measured in trades number) when the current trade is a trade-through. A similar graph is obtained when the lag is measured in physical time. We obtain a mean of 30 trades (respectively 113 seconds) to wait for the next trade-through in the unconditional case and a mean of 23 trades (respectively 91 seconds) when already on a trade-through.

To verify that this difference is statistically significant, we use a Welch two samples t-test to compare the means of the distributions in the conditional and unconditional cases and obtain an equality test p-value of less than 2.2×10^{-16} in both cases. So, trade-throughs are both more likely followed by trade-throughs (in trades number), and more closely in time followed by trade-throughs.

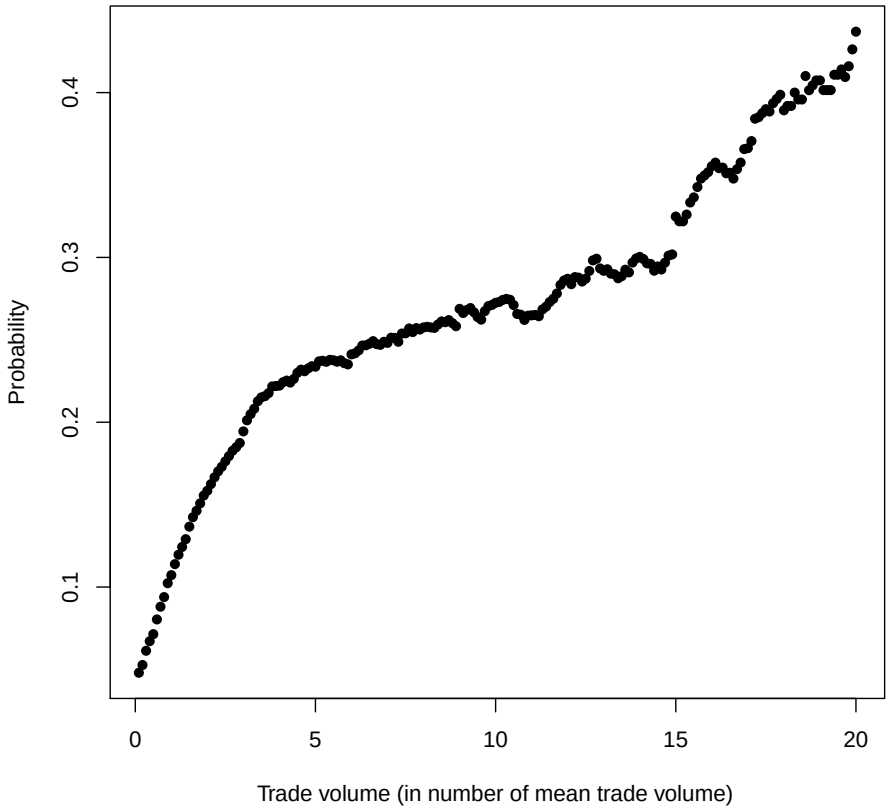

Figure 3. Conditional probability for a trade to be a trade-through conditioned on his volume for BNP Paribas stock (in linear-logarithmic scale)

3.5. Intraday distribution

Figure 5 shows the intraday distributions of timestamps[1] of trades in the unconditional case (all trades) and the conditional case (trade-throughs) for the US equity future E-mini S&P500. Comparing the distribution for 2nd-limit trade-throughs with the unconditional one shows that the U-shaped central part is less important, and that the relative size of the peaks drastically increases. Those peaks are present at very specific hours, observed for both US and European equity futures, as will be detailed later. Note that if we restrict the study to 3rd-limit trade-throughs, the U-shape part is almost completely removed from the distribution.

Using a two-sample Kolmogorov-Smirnov test to compare the unconditional and conditional distributions of the timestamps of trade-throughs yields a p-value less than 2.2×10^{-16}, which indicates that both distributions are statistically different.

The peaks of the intraday distribution of trade-throughs timestamps are more pronounced at

[1]All timestamps presented in this article are referenced in the time-reference of Europe/Paris = CET = UTC/GMT + 1h.

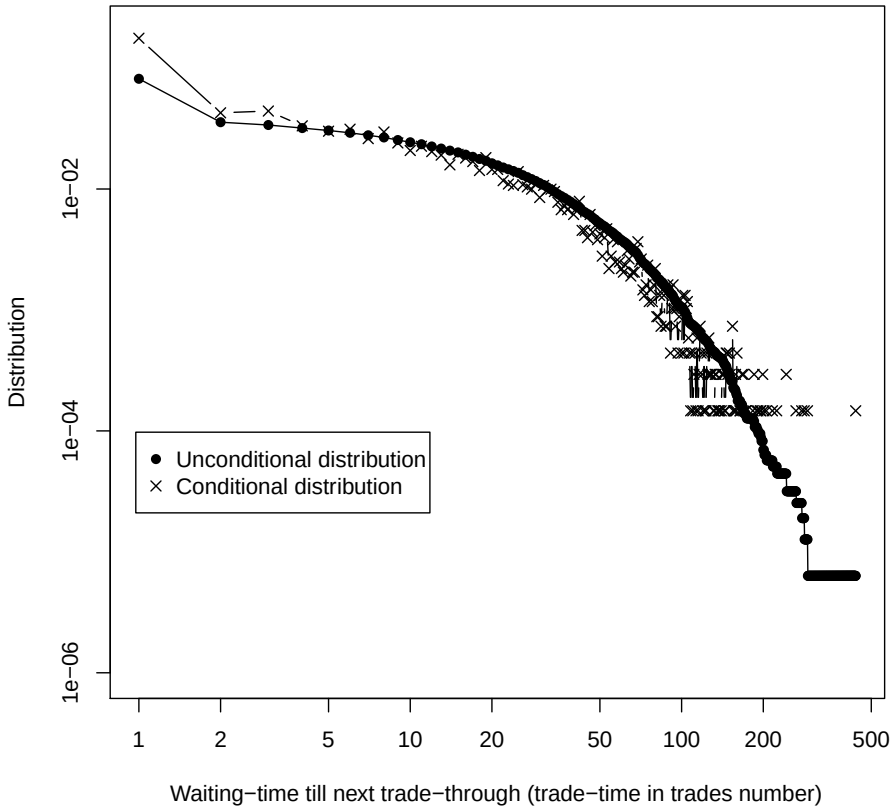

Figure 4. Distribution of waiting-time (in trade-time) until next trade-through for BNP Paribas stock (in double logarithmic scale)

specific hours of the day :

- 07.50 : Eurex trading phase beginning (FESX, FDAX).
- 09.00 : Euronext trading phase beginning (FTSE).
- 14.30 : US major macro news releases (Jobless claims, Employment situation, International trade or GDP, for example). CME open-outcry trading phase beginning (major equity index futures).[1]
- 15.30 : NYSE regular trading phase beginning.
- 16.00 : US major macro news releases (ISM Manufacturing Index, Philadelphia Fed survey or New Homes Sales, for example).
- 17.30 : End of the calculation of the DAX index using Xetra electronic trading system.
- 22.00 : Eurex trading phase end (FESX, FDAX). Euronext trading phase end (FTSE).

Similar results are obtained for European Equity futures and French stocks, but with smaller

[1]The possible links between news and major events in stocks have already been studied, for example by Joulin *et al.* (2008)

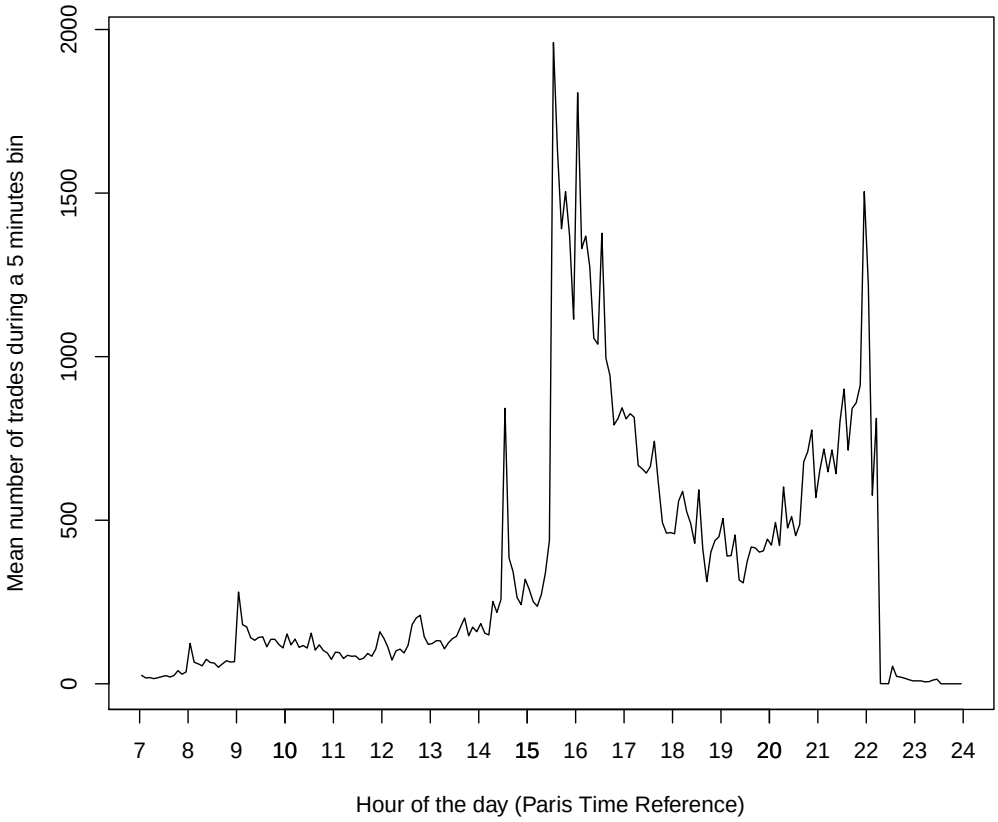

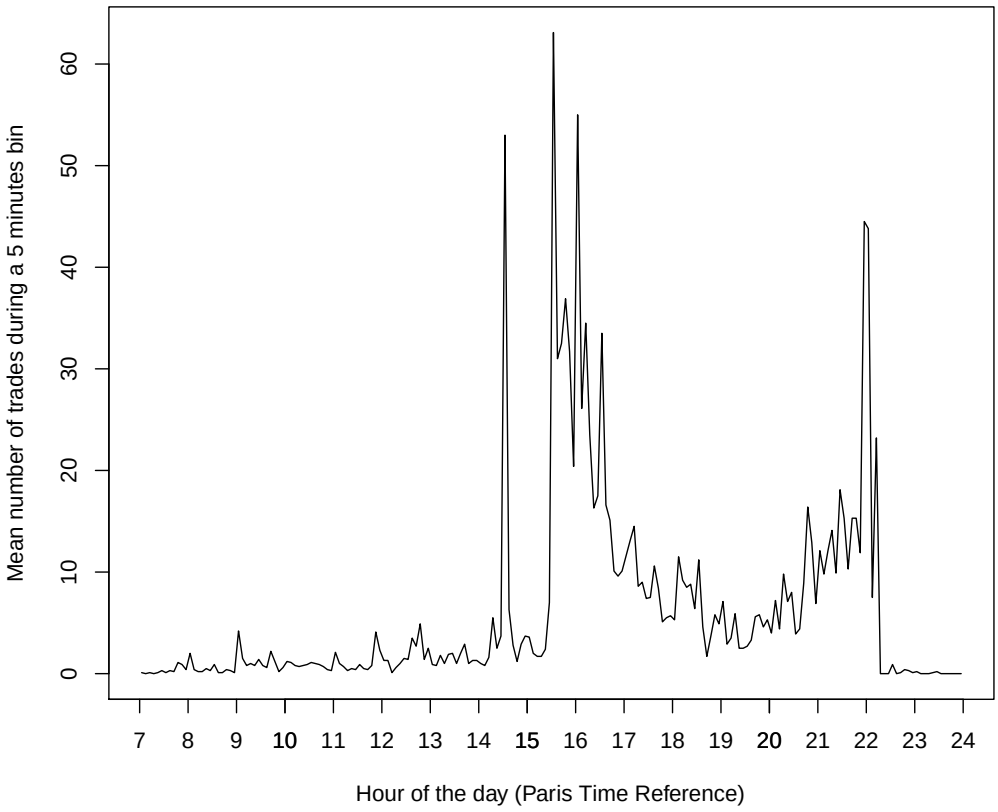

Figure 5. Trades (above) and trade-throughs (below) intraday histograms for US equity index future E-mini S&P500

peaks for French stocks. This suggests a weaker dependence of French stocks on major macroeconomic news, compared with US and EU equity futures.

4. Market impact

In double auction limit order books, the price is completely determined by the liquidity available in the order book and the way liquidity is consumed by market orders. An interesting and important question is to assess the market impact, measured as the correlation between the sign of an incoming market order and the subsequent price changes. An empirical measure of market impact is given in Bouchaud *et al.* (2004) by the impact function R as a function of the lag l:

$$R(l) = \langle (m_{n+l} - m_n).\epsilon_n \rangle$$

where ϵ_n is the sign of the n-th trade (1 if the trade is on the ask side, -1 if it is on the bid side) and m_n is the midprice before the n-th trade.

Eisler *et al.* (2009) studied the market impact of 14 of the most liquid stocks traded on the NASDAQ in 2008. They obtained a response function that increases for lags up to 100 events and then, depending on the tick size, remains roughly constant for large tick stocks, or slowly decreases to a non-zero value for small tick stocks. In this section, we adapt the methodology in Eisler *et al.* (2009) to measure the market impact of trade-throughs. One important point has to be taken into account: by definition, there is a mechanical, instantaneous price change after a trade-through that has to be removed to fairly compare the price impact of a trade-through with that of any trade. To this end, we start counting the lags one instant [1] after the considered trade-through. That way, we make sure that the reference level with respect to which the returns are calculated is the price after the trade-through, and not before it.

4.1. Empirical response function

In Figure 6, are plotted the response functions in the unconditional case (all trades are considered) and in the conditional case (only trade-throughs), for lags smaller than 15 minutes. One can see that the two response functions are similar: they first sharply increase from 0 to a maximum value (equal to 1 tick for trade-throughs and to 0.8 tick for all trades) reached after approximately 10 seconds, and then slightly decrease to half of its maximum value for trade-throughs, and 75% of its maximum value in the unconditional case as the time lag goes from 10 seconds to 15 minutes. Also note that the response function of trade-throughs is larger than the unconditional one for lags less than one minute, whereas both response functions become approximately equal for lags ranging from 1 minute to 15 minutes.

The global response function clearly measures the correlation between the sign of a trade and that of the next return. One can see that for lags less than 1 minute, the sign of a trade-through is a better indicator of the sign of the next midprice return than the sign of any trade. In an economic perspective where information is displayed in the market through the combination of events affecting the limit order book, one can say that trade-throughs have a

[1] In the data, we use a time lag of one millisecond.

higher informational content than usual trades.

Comparing our results with those in Eisler *et al.* (2009), we first notice a difference in the order of magnitude of the price impact function. After removing the instantaneous price impact in their graphs, their response functions take values ranging from 0 to 0.4 tick for large tick stocks and from 0 to 0.8 tick for small tick stocks, whereas in our example, the response function is more pronounced and takes values ranging from 0 to 1 tick. Apart from the instantaneous price impact, they obtain very similar response functions for market orders that do not change the best price and for those that change the best price. The only difference between both cases is, in the case of small tick stocks and for lags larger than 2000 events, they find a constant impact for market orders that change the price, and a response function that continues to decay for market that do not change the price. In the study we perform on trades-through, the response functions for small and large tick stocks both decay for lag values of 10 seconds to 15 minutes. Given the number of events per day, of the order of 70 000 events per day on the French stock BNP Paribas, this is most likely caused by the fact that we do not reach the regime where Eisler *et al.* (2009) began to measure a constant impact function, as we do not concatenate successive days when performing our statistics.

4.2. Response function conditioned on trade volume

We have seen in Sections 3.2 and 3.3 that trade-throughs tend to be trades with volume bigger than the mean trade volume. It is also clear that trade volume influences market impact: the larger the trade size, the greater the impact. The results shown on Figure 6, where we compared global response functions for trade-throughs and all trades, may arguably be explained by the fact that trade-throughs generally have bigger volume. To rule out this simplistic explanation, we include the volume of trade-throughs in the response function and study the difference in market impact.

Following Bouchaud *et al.* (2004), we use a generalized version of market impact to take the trade volume into consideration:

$$R(l, V) = \langle (m_{n+l} - m_n).\epsilon_n \rangle \,|_{V_n = V}$$

where ϵ_n is the sign of the n-th trade, m_n is the midprice before the n-th trade and V_n is the volume of the n-th trade.

Trades are divided into four volume categories related to the mean trade volume MTV: $[0, \frac{MTV}{2}], [\frac{MTV}{2}, MTV], [MTV, 2 \times MTV], [2 \times MTV, +\infty]$. Figure 7 shows that each response function of trade-throughs is still bigger than the unconditional one of the same trade volume category. This property holds true for all lags for the first trade volume category $[0, \frac{MTV}{2}]$ and for lags ranging from 0 to 10-20 seconds for the other trade volume categories.

One must therefore conclude that, independently of the trade volume, the sign of trade-throughs is a better indicator of the sign of the next return than the sign of an unconditional trade, at least over timescales ranging from 0 to 10 seconds. This confirms what may be termed as the higher informational content of trade-throughs as opposed to that of usual trades[1].

[1]Note that similar empirical results are obtained using other time definitions, e.g. tick time, or trade volume time where the lag is measured by the fraction of trading volume with respect to the daily trading volume.

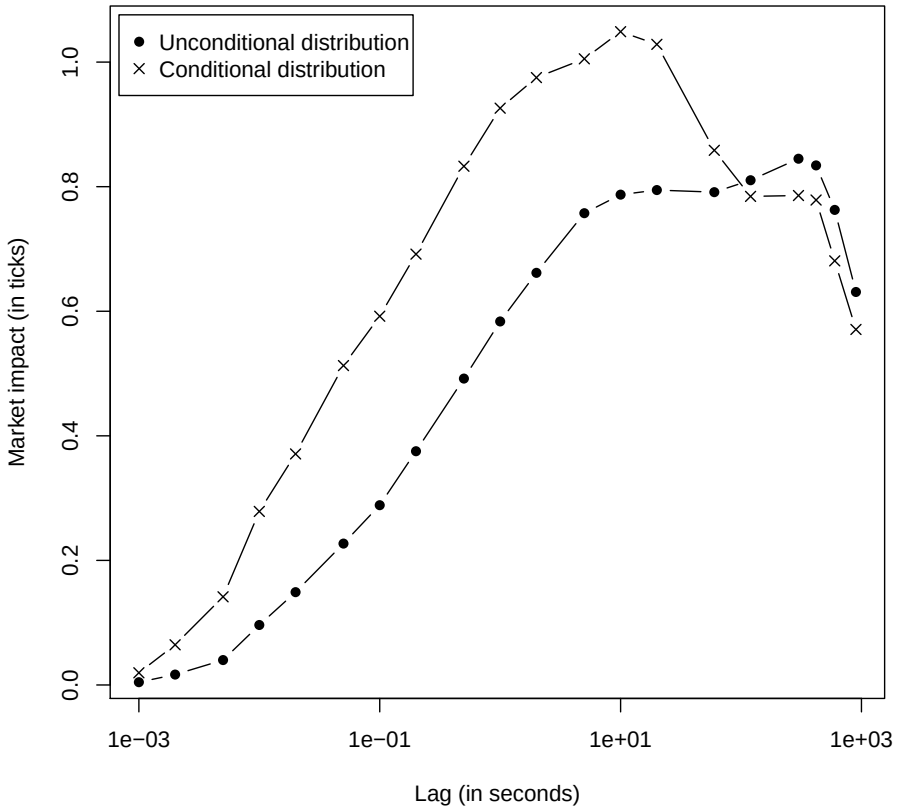

Figure 6. Global response functions (all trades vs trade-throughs) for BNP Paribas stock (in linear-logarithmic scale)

5. Spread relaxation

As trade-throughs are rather rare and informationally rich events, it is natural to consider a trade-through as an excitation of the limit order book. The results of the previous sections essentially addressed the relaxation of the mid-price, and we now focus on the behaviour of the spread. Hence, we study the spread of the limit order book in order to measure:

- the spread level in the excited state (on a trade-through)
- the evolution of the spread after the trade-through.

Figure 8 shows the behaviour of the excess spread for the French stock BNP Paribas after all trades and after trade-throughs. The excess spread is defined as the difference between the value of the spread after a trade and its value just before. The excited value (in ticks) of the excess-spread is approximately 0.5 after a trade-through, and 0.25 after a standard trade. Since trade-throughs instantaneously increase the spread by one tick, this result indicates that approximately half of the relaxation takes place during the first lag interval, namely, 5 seconds.

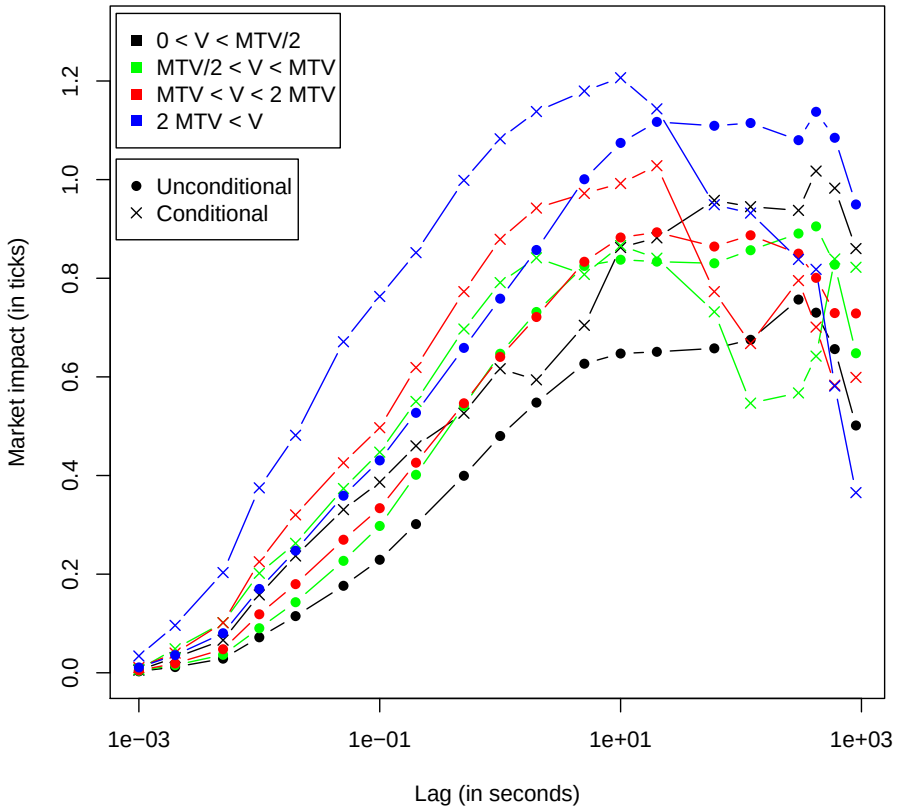

Figure 7. Global response functions (all trades vs trade-throughs) for BNP Paribas stock with the condition on the initial trade volume with respect to the mean trade volume MTV (in linear-logarithmic scale)

The unconditional excess spread decay seems to decay almost linearly for lags smaller than 100 seconds. For longer lags, one must bear in mind that the scarcity of available data where two successive trade-throughs are separated by a large lag makes a large portion of this unconditional curve not statistically reliable. It is plotted it here mostly for indicative purpose, and is mostly reliable for the very first lags.
For the sake of completeness, one can fit a power-law decay of the excess spread in physical time with an exponent close to 0.25, as shown on Figure 9.

At this stage, a comment is in order: Section 3.5 shows that there is approximately one trade-through per minute for BNP Paribas stock. However, there are periods during the day when more events occur, and the maximum number of trade-throughs reaches a frequency of 1 event every 25 seconds. Therefore, the excess-spread relaxation is in fact statistically reliable for a period of the order of 25 seconds. For lags larger than 25 seconds, there are significantly less data used in the conditional statistics and one should be very careful drawing any conclusion from that part of the graph.

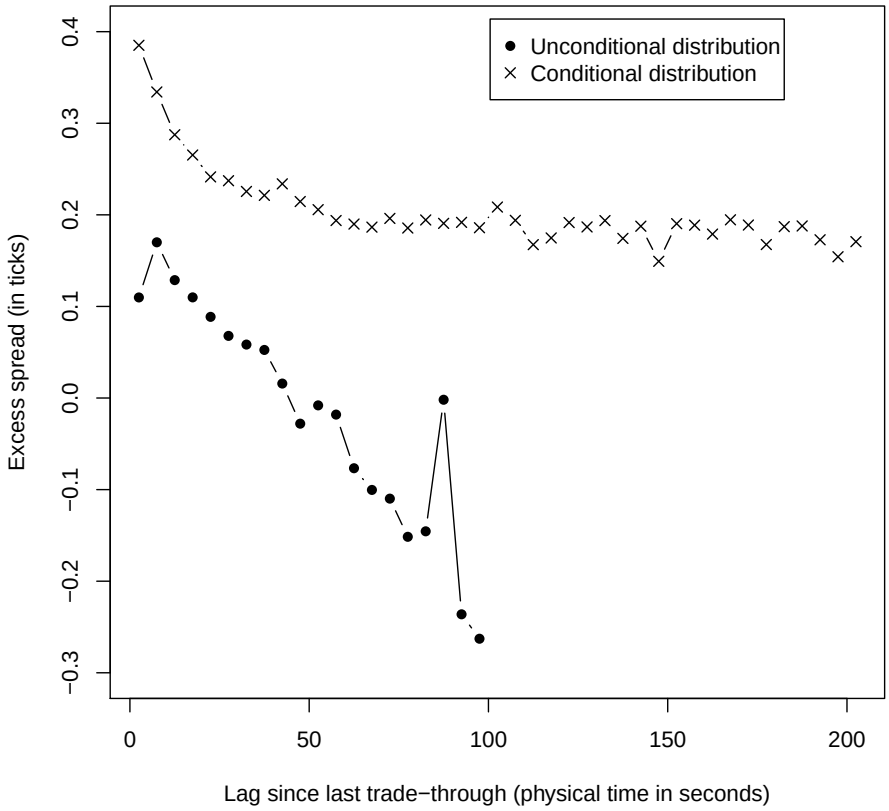

Figure 8. Excess-spread decay after trades for BNP Paribas stock (in linear-logarithmic scale)

Also note that we consider physical time in this section. Tick time could be relevant, as it incorporates all the changes in the order book. However, it seems clear that trade time, especially trade-through time, is not the correct clock, since most of the spread dynamics takes place before a next trade occur, as the order book replenishes itself.

Ponzi *et al.* (2009) study a similar problem in the relaxation of the spread. Conditioned on a move of the spread, they measure a relaxation and obtain a power-law behavior of the excess spread in trade time with exponent between 0.4 and 0.5). They provide no explanation for this empirical observation.

Another example is given by the work of Toth *et al.* (2009) where relaxation after large price moves is studied. They show a power-law relaxation of the excess bid-ask spread in physical time, with exponent 0.38.

In both cases, the methodologies and results are similar to ours, in that they provide evidence of a slow relaxation of the excess-spread, with power-law fits showing exponents of the same order of magnitude.

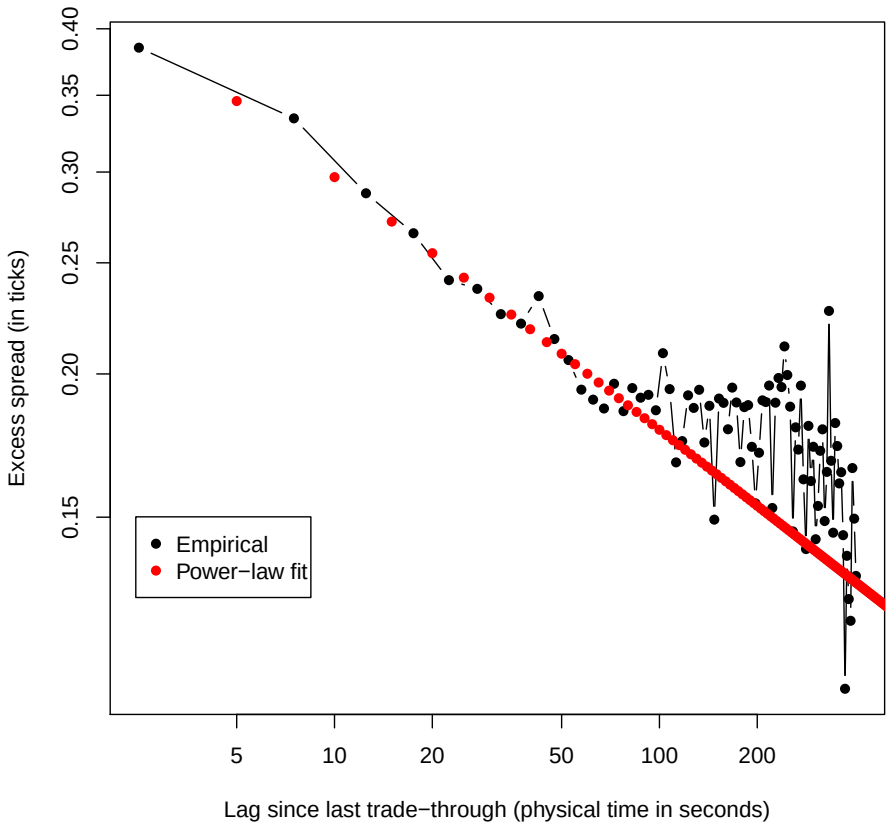

Figure 9. Power-law decay of excess-spread decay after trade-throughs for BNP Paribas stock (in double logarithmic scale)

6. Lead-lag estimation from trade-throughs time-series

Finding lead-lag relationships between two different assets is an important statistical problem in finance. As said in the Introduction, there are recent references in the literature (see Abergel and Huth (2011) and Hoffmann and Rosenbaum (2011)) aiming at finding empirical results and theoretical methods to evaluate lead-lag relationships based on the lagged cross-correlation of the returns time series. Our approach is different, as we directly address the following question: which asset moves first ? In other words, when a trade-through occurs on Asset 1, is it followed or is it preceded by a trade-through on Asset 2 ? Thanks to the relative rarity of trade-throughs, we are able to answer this question in a statistically significant way, as will now be described in details.

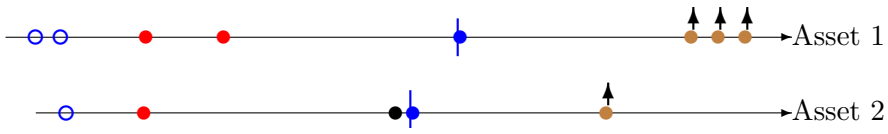

Figure 10. trade-throughs of Asset 1 are put in relationship with those of Asset 2

6.1. The estimation technique

Assume that we have two grids representing the timestamps of trade-throughs for two different assets, and that we want to get an empirical distribution of the lead-lag parameter between the two assets. Our method is quite simple: we connect every timestamp on one grid with its nearest neighbour on the other grid. At the end of this process, every trade-through on Asset 1 is linked with the closest trade-through on Asset 2. Figure 10 shows how different timestamps are connected pairwise.

The empirical distribution of the lead-lag parameter between the two assets is then obtained by calculating the difference between two connected timestamps and plotting this distribution.

This method can be generalized to two groups of more than one asset by first merging the timestamps of trade-throughs in each group. For example, in the study of the lead-lag between EU and US equity markets, we first build two grids, one for the US equity market and one for the EU equity market, by merging the timestamps of the trade-throughs of E-mini S&P500, Nasdaq E-mini and Dow Jones E-mini futures (for the US grid) and those of Eurostoxx, DAX and Footsie futures (for the EU grid).

In order to better understand this estimation technique, and also provide numerical evidence that it is sound, we test it with a very basic model for the arrival times of trade-throughs as two independent Poisson processes with different intensities. In such a model, there is no lead-lag relationship between the two assets. On Figure 11, are plotted the two empirical lead-lag distributions using either one of the assets as a reference grid. Clearly, the two distributions are different, but they are both perfectly symmetric with zero mean.

An important feature of this method is that it is purely empirical, totally model-free, and returns a full distribution of lead-lag parameter, not only one value that maximizes some contrast criterion. Again, let us insist on the fact that only the relative rarity of trade-throughs enables us to use this simple method: should the average lead-lag parameter be of the order of the average frequency of trade-throughs, no conclusion could be drawn.

6.2. Empirical results

In this section, we apply the methodology just introduced to empirical lead-lag estimates to the EU and US equity futures. Results are displayed on Figure 12, where two useful statistics are plotted: the empirical distribution of the lead-lag parameter, and the positive and negative cumulative distributions of the lead-lag parameter. On the left side of each graph, the EU futures lead the US futures and on the right side, the US lead Europe.

Define the lag l between the timestamps TS of the two stocks as $l = TS_{EU} - TS_{US}$. We have to

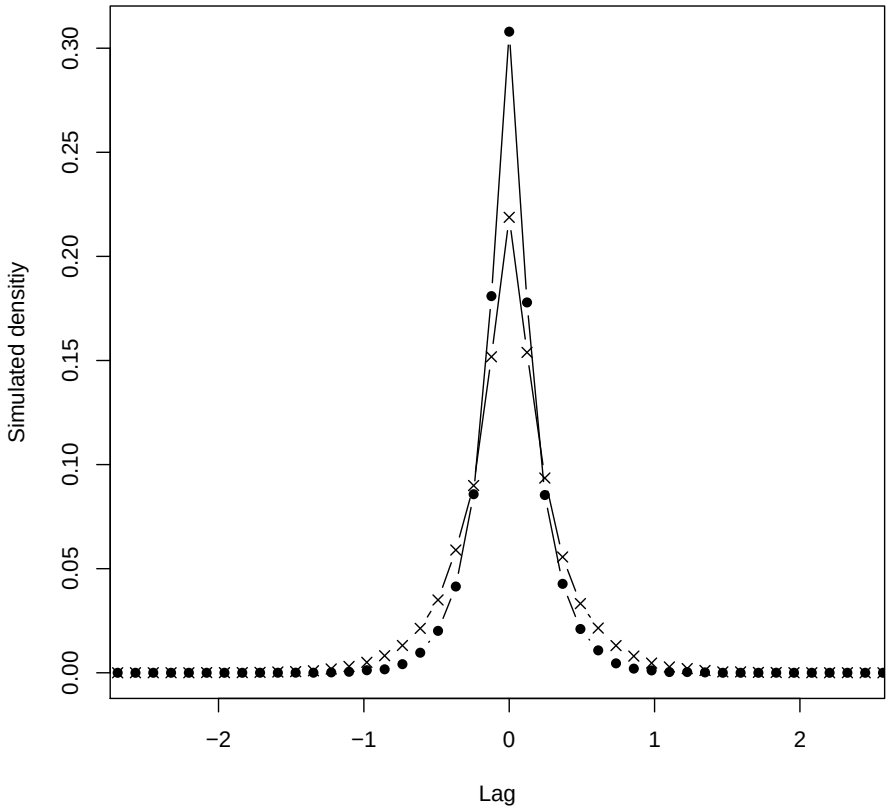

Figure 11. An example of lead-lag parameter distribution for two simulated Poisson processes with different intensities (one curve for each time grid reference)

compare $P_- = \mathbf{P}(l < 0) \equiv \mathbf{P}(\text{EU leads US})$ and $P_+ = \mathbf{P}(l > 0) \equiv \mathbf{P}(\text{US lead EU})$. Empirically, there holds that $P_+ > P_-$. Let us now check that this difference is statistically significant, by comparing it to its standard deviation σ. Since

$$\sigma^2 = \mathbf{Var}(P_+ - P_-) = \mathbf{Var}(2P_+ - 1) = 4\mathbf{Var}(P_+) = \frac{4P_+P_-}{N_{data}}, \tag{1}$$

the empirical values $P+ \approx 0.5126$, $P- \approx 0.4874$, $N_{data} = 118207$ yield $P_+ - P_- \approx 0.025$ and $\sigma \approx 0.0029$, and therefore $P_+ - P_- \approx 8\sigma$. One can then conclude that the empirical lead of US futures on EU futures is statistically significant at an 8 standard deviations level on the time period considered.

We also perform a similar analysis on French stocks, choosing the pair Société Générale and BNP Paribas. Defining similarly the lag l between the timestamps TS of the two stocks as $l = TS_{SOGN} - TS_{BNP}$, we obtain $P_- = \mathbf{P}(l < 0) \equiv \mathbf{P}(\text{SOGN leads BNP})$ and $P_+ = \mathbf{P}(l > 0) \equiv \mathbf{P}(\text{BNP leads SOGN})$. Clearly, $P_- > P_+$. The empirical values $P+ \approx 0.489$,

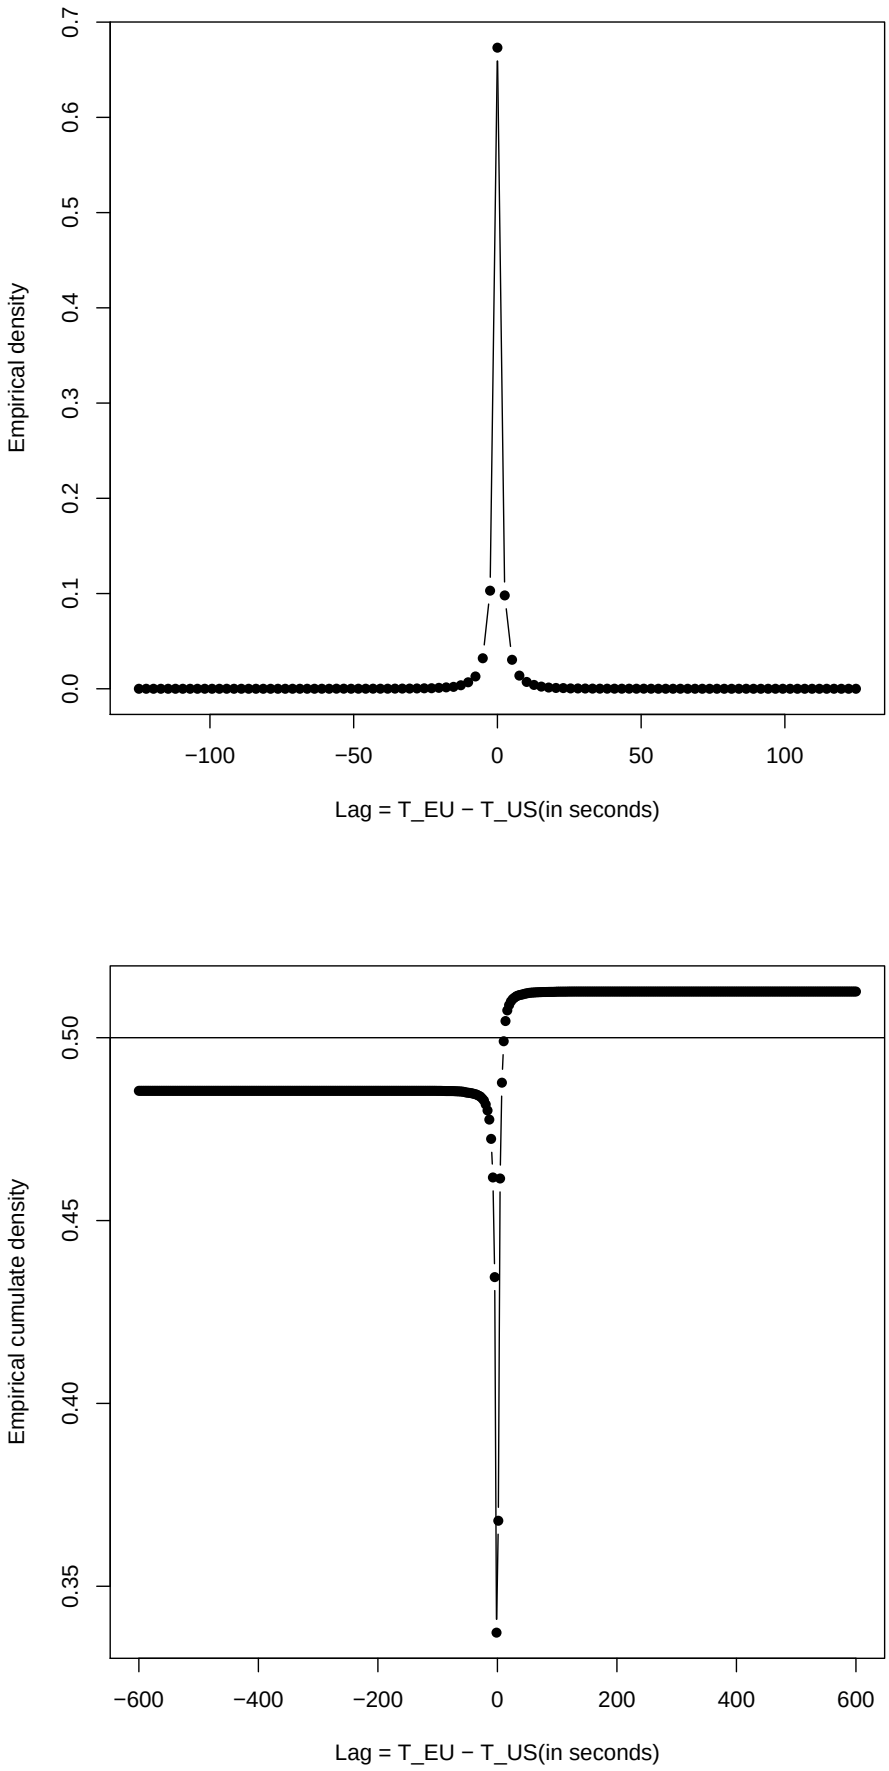

Figure 12. Lead-lag parameter empirical distribution (above) and positive and negative cumulative distributions (below) of 2nd-limit trade-throughs for two groups of European and American futures

$P- \approx 0.511$, $N_{data} = 6816$ yield $P_- - P_+ \approx 0.022$ and $\sigma \approx 0.012$, which means that the lead of Société Générale on BNP Paribas is statistically significant at approximately a 2 standard deviations level.

7. Conclusion and further research

In this paper, we perform an extensive study of multiple-limits trades, or trade-throughs. Various statistics are provided measuring their liquidity, volumes, arrival times, clustering and spread relaxation. Evidence is provided regarding important peaks in the arrival time distribution at 2.5 pm and 4pm (Paris time reference), the time of the day when major macro-economic news are released. Market-impact of trade-throughs is studied, and is demonstrated to be larger than that of other trades, a fact that we synthetize by saying that trade-throughs have a higher informational content. We also introduce a new methodology to assess lad-lag relationships, and apply it to pairs of financial markets or assets.
Finally, we want to mention a recent related work, Muni-Toke and Pomponio (2011), where the mathematical modelling of trade-throughs using Hawkes processes is explored.

Acknowledgments

The authors would like to thank the members of BNP Paribas Equities & Derivatives Quantitative R&D team for fruitful discussions, and especially Boris LEBLANC and Stéphane TYC. The authors also want to thank the anonymous referees for their very thorough reading of the first version of the manuscript, leading to major improvements at several places.

References

Abergel, F. and Huth, N., High frequency lead/lag relationship-empirical facts, *Arxiv preprint. arXiv:1111.7103v1,* 2011.

Almgren, R., Thum, C., Hauptmann, E. and Li, H., Direct estimation of equity market impact, *working paper* 2005.

Bouchaud, J.-P., Farmer, J.D. and Lillo, F., How markets slowly digest changes in supply and demand, *Handbook of financial markets: dynamics and evolution*, North-Holland, Elsevier 2009, 57-160.

Bouchaud, J.-P., Gefen, Y., Potters, M. and Wyart, M., Fluctuations and response in financial markets : the subtle nature of 'random' price changes, *Quantitative Finance,* 2004, **4(2)**, 176-190.

Bouchaud, J.-P. and Potters, M., Theory of Financial Risk and Derivative Pricing: from statistical physics to risk management, *Cambridge university press* 2003.

Cont, R., Kukanov, A. and Stoikov S., The price impact of order book events, *preprint available at http://arxiv.org/pdf/1011.6402.pdf*, 2011.

Eisler, Z. , Bouchaud, J-P. and Kockelkoren, J., The price impact of order book events: market orders, limit orders and cancellations, *Quantitative Finance, to appear* 2011, *preprint available at http://arxiv.org/pdf/0904.0900v3.pdf.*

Farmer, J. D., Gillemot, L., Lillo, F., Mike, S. and Sen A., What Really Causes Large Price Changes?, *Quantitative Finance,* 2004, **4**, 383-397.

Gerig, A., A Theory for Market Impact: How Order Flow Affects Stock Price, *PhD thesis, University of Illinois* 2007.

Hautsch, N., and Huang, R., The market impact of a limit order, Sfb 649 discussion papers, Humboldt University, Berlin, Germany, 2009.

Hoffmann, M., Rosenbaum, M. and Yoshida, N., Estimation of the lead-lag effect from nonsynchronous data, *Bernoulli, to appear* 2011.

Joulin, A., Lefevre, A., Grunberg, D. and Bouchaud, J.-P., Stock price jumps: news and volume play a minor role, 2008. *Wilmott Magazine,* 2008, **vol. Sep/Oct**, 1-7.

LeBaron, R., Yamamoto, R, Long-memory in an order-driven market, *Physica A : Statistical Mechanics and its Applications* 2007, **383(1)** 85 -89.

Lillo, F. and Farmer, J.D., The Long Memory of the Efficient Market, *Studies in Nonlinear Dynamics & Econometrics,* 2004, **8(3)**.

Lillo, F., Mike, S. and Farmer, J.D., Theory for long memory in supply and demand, *Physical Review,* 2005, **E 71, 066122**.

Muni-Toke, I. and Pomponio, F., Modelling trade-throughs in a limit order book using Hawkes processes, *Economics Discussion Papers*, 2011, **No 2011-32**, *preprint available at http://www.economics-ejournal.org/economics/discussionpapers/2011-32*

Ponzi, A., Lillo, F. and Mantegna, R. N., Market reaction to a bid-ask spread change: A power-law relaxation dynamics, *Physical Review,* 2009, **E 80, 016112** .

Potters, M. and Bouchaud, J.-P., More statistical properties of the stock order books and price impact, *Physica A : Statistical Mechanics and its Applications* 2003, **324(1-2)**, 133-140.

Toth, B. , Kertsz, J. and Farmer, J.D., Studies of the limit order book around large price changes, *Eur. Phys. J.,* 2009, **B 71**, 499–510.

Weber, P. and Rosenow, B., Order book approach to price impact, *Quantitative Finance,* 2005, **5(4)**, 357-364.

Weber, P. and Rosenow, B., Large stock price changes: volume or liquidity?, *Quantitative Finance,* 2006, **6(1)**, 7-14.