Significance and impact of high-frequency trading in the German capital market

Over the past ten years, the significance of algorithm-based trading strategies has grown considerably in international marketplaces, especially in Europe. This has accordingly heightened interest on the part of central banks and regulators in the potential implications of high-frequency trading (HFT) on market stability and market integrity.

However, the market impact of HFT to date has been all but impossible to measure, especially owing to the paucity of available data. In order to provide a stronger basis for a discussion based on facts, this report presents the first comprehensive empirical studies of HFT in the German capital market.

The results illustrate the fact that it would be inappropriate to make a sweeping judgment of the impact of HFT on the financial markets as they depend heavily on the strategies and market phases under observation, which are also, in some cases, quite varied. This means that, in a calm market setting, HFT traders make a considerable contribution to liquidity. However, during periods of high market volatility, the studies show that HFT market makers temporarily reduce their liquidity provision in both Bund and DAX futures. In addition, in times of relatively highly volatile markets, some HFT participants are particularly active and can thus contribute to trend-amplifying price movements. The results can also stimulate the regulatory debate on HFT.

Introduction

HFT accounts for key share of exchange trading

Since the middle of the last decade, investors, exchange operators and regulators have been witnessing a remarkable transition in the ways in which securities and derivative financial instruments are traded in the financial markets. Computer algorithms have been making increasing inroads into exchange trading activities, for instance. Human decisions on which securities to buy or sell and at what prices are being replaced with increasing frequency by specially designed algorithms which are capable of analysing large quantities of data and initiating hundreds of orders in a fraction of a second. This trend has been accompanied by growing competition among stock exchanges, new regulatory measures (MiFID in Europe, Reg NMS in the United States) and the appearance of new marketplaces which, in particular, favour algorithm-based trading.[1] These various factors, acting in concert, constitute a driver of structural change in the financial system.

High-frequency trading (HFT), as a sub-category of general computerised trading, is playing a key role in this transformative process. It now accounts for nearly 50% of trading activity in the most highly liquid segments of the US and European markets.[2] HFT uses new technological infrastructures and algorithms in order to profit from high trading speeds based on a variety of different strategies. Although the use of speed-based advantages in exchanges is not an unprecedented approach by any means, the advent of HFT has made speed a particularly important factor.

Recent events on stock exchanges have pushed HFT to forefront of public debate

At the beginning of the current decade, HFT drew the attention of the public and regulators owing to a series of events on stock exchanges[3] which were characterised primarily by price swings which were rapid and, in most cases, could not be explained by the fundamentals. Many observers saw these events as having been triggered by the activity of HFT algorithms. Recently, such "flash events" have also been increasingly occurring in the markets for sovereign bonds, traditionally regarded as highly liquid and less volatile. Examples include the US Treasury "flash rally" in October 2014 and the Bund tantrum[4] in early 2015.[5] These extreme events triggered debates on financial markets' ability to withstand shocks – in other words, their resilience – and on the impact of new market players, such as HFT.

Impact of HFT on aspects of market quality highly interesting …

Given that HFT accounts for a high share of trading activity, together with the fact that opinion on it among investors is sharply divided and that there is a general interest in new forms of capital market intermediation, regulators and central banks have been increasingly addressing the issue of HFT for some years now. They are looking not only at market integrity issues but also first and foremost at the impact HFT has on the ability of markets to function. To the extent to which HFT impacts

1 The Markets in Financial Instruments Directive (MiFID) is an EU directive designed to harmonise financial markets in the single European market. Regulation (Reg) NMS (short for "national market system") is a US financial directive which is intended to enable all investors to place orders at the best nationwide rate.
2 It is very difficult to precisely quantify the HFT share of the volume of a specific traded asset. For the US equity markets, a study by the TABB Group for 2012 ("US Equities Market 2012: Mid-year review") finds that the share is around 50% (as early as 2005, this share had already stood at roughly 30%). According to a more recent ESMA study in 2014, the share of equities trading on European stock exchanges accounted for by HFT, depending on the underlying definition applied, ranges from 24% to 43% of the equity trading volume and 58% to 76% of the overall number of orders.
3 A "flash event" is a sudden occurrence of very strong and rapid price volatility, for which a fundamental cause is often absent and the path of which is generally reversible. Examples include the "flash crash" of the US stock market on 6 May 2010 and a whole series of smaller, but no less dynamic, events in a variety of asset classes and exchanges. They have also marked the beginning of a period of intensive study of the topic by researchers and regulators since 2010.
4 The Bund tantrum, which occurred on 7 May 2015, is seen as the peak of a period of strong price volatility in German government bonds (Bunds) which lasted from late April to early June 2015.
5 Central banks and regulatory authorities, in particular, responded by publishing studies in which they examined the history and causes of these events. See US Department of the Treasury (2015), Joint Staff Report, the US Treasury market on October 15th 2014; R Riordan and A Schrimpf (2015), Volatility and evaporating liquidity during the bund tantrum, BIS Quarterly Review, September 2015; BIS (2016), Electronic trading in fixed income markets, Study by the Market Committee.

on various facets of market quality, such as liquidity, volatility and price efficiency, it can also be expected to impact on financial stability.

… yet hardly any empirical studies for German capital market to date

Owing primarily to a paucity of available data, empirical knowledge of the actual extent of HFT activity and its impact is currently limited. Studies on European capital markets and for non-stock-market segments – including Germany – are particularly few and far between. Another factor which makes it difficult to come to a definitive judgement on the impact of HFT is that HFT is simply a portmanteau for a wide variety of strategies. This heterogeneity makes it a laborious endeavour to identify a unique impact of HFT on market quality.[6]

Focal points of study: liquidity provision, price efficiency and price volatility

With all that in mind, the present article was written with a view to helping the reader to obtain a deeper understanding of the importance of HFT activities and their impact on the German capital markets. One of the aims is to address whether the rising velocity of financial market activities is having an overall positive impact on the capital markets.[7] However, the article is chiefly about the impact of HFT on liquidity provision during various market phases and the role of HFT in information processing and its contribution in periods of temporary high volatility. Some studies look at the market for DAX and Bund future contracts at the micro-second level, the behaviour of HFT traders when key macroeconomic data are published and the role of HFT as a market maker in various market settings. In addition, the DAX future order book is extensively reconstructed and evaluated and market data analysed for potentially conspicuous strategies. One particular advantage of these studies over previous research is that current and extensive data from Eurex, in which each individual HFT order is labelled, are used.

HFT strategies and behaviours in different market phases

Definition of HFT pursuant to MiFID II

MiFID II defines HFT as a subcategory of algorithmic trading. The key characteristics of HFT include[8]

– a technical infrastructure designed to minimise the time lag for order executions and

– which uses the opportunity, for a fee, to station its trading computers in close proximity to the exchange’s servers (co-location) and/or very high-speed electronic access for system-determination of the initiation, generation, forwarding (“routing”) and execution of individual trades or orders without any human intervention

– with high intraday message rates of orders, quotes or deletions.

Overview of HFT strategies

HFT comprises numerous different strategies …

HFT is a portmanteau term that covers a large number of different strategies. HFT strategies can be broken down into the following categories, depending on their objective: statistical arbitrage, directional strategies and passive market making.[9] Moreover, there are what are known as structural strategies which are based on exploiting structural weaknesses in the market infrastructure or among individual market players. Other strategies are also the topic of critical debate in the public sphere and among academics.[10]

6 See US Department of the Treasury (2015), op cit.
7 For more information on this fundamental category, see, for example, J Weidmann (2014), The macroeconomic importance of capital markets, speech delivered on 22 May 2014.
8 For more information, see MiFID II, Article 17 (1) and (2).
9 See Securities Exchange Commission (2014), Equity market structure literature review part II: high frequency trading, Working Paper.
10 These include quote stuffing (see box on pp 55-58), momentum ignition and strategies whose sole purpose is to detect movements in the order book and identify liquidity (eg sniping).

… such as statistical arbitrage …

Statistical arbitrage is designed to exploit price inefficiencies between related products or markets. For example, such a strategy takes advantage of temporary price discrepancies between highly correlated products, such as an exchange-traded fund (ETF) and the underlying basket of individual securities.

… directional strategies and passive market making as the best-known categories

In directional strategies, HFT traders take unhedged positions based on anticipated price changes and mostly trade in the direction of the short-term price trend. One type of directional strategy is "news trading", in which the objective is to respond to new public information as fast as possible in order to generate fast profits (see the box on pages 47 to 49).

Passive HFT traders, acting as market makers, provide liquidity and continuously update their bid-ask spreads to reflect the market situation. Their primary sources of income include, alongside the bid-ask spread, fees for liquidity provision services offered by some trading platforms, primarily in the United States.[11]

Impact on market stability depends on type of strategy and market setting

Some of the strategies being pursued by HFT traders are suspected of amplifying or even causing market turmoil, depending on the market setting and aggressiveness of their implementation. Therefore, a closer look should be taken not only at market integrity but also primarily at the impact on stability. While the statistical arbitrage is not an issue here, as it only causes the prices of various products on one platform or the same products on different platforms to adjust more quickly, directional strategies, in particular, can contribute to stronger price movements. Whether or not they also cause major price swings depends not least on the liquidity provided during the respective market phase. Therefore, the focus of the empirical studies in the next sections will be, in particular, on directional and passive market making strategies.

Behaviour of HFT participants in different market phases

The following section will study the general trading behaviour of HFT participants and slower traders before going into detail on specific HFT strategies. It will also study whether the behaviour of market players differs in periods of high and low volatility.[12] On the basis of the implied volatility index on the DAX (VDAX), a trading week of heightened volatility (March 2014), as well as a week of low volatility (June 2014) are identified. The underlying Eurex data are described in the box on page 42.

Active trading orders consume liquidity, passive orders provide liquidity

For a more detailed analysis of HFT participants' trading behaviour, orders initiated by market participants will be divided into active and passive orders.[13] Market orders or limit orders which are executed immediately[14] (liquidity-consuming) are regarded as active orders. By contrast, limit orders which are not executed immediately and which transmit liquidity to the order book (liquidity-providing) are regarded as passive orders. In the following, market participants which issue an active (passive) order at a given point in time are referred to as active (passive) HFT participants.

11 See SEC (2010), Concept release on equity market structure, Working Paper. These "maker-taker" pricing systems have been used increasingly on electronic marketplaces since the late 1990s, including the NYSE and the NASDAQ.
12 For details regarding the method, see K Schlepper, High-frequency trading in the Bund futures market, Deutsche Bundesbank Discussion Paper No 15/2016; as well as J Breckenfelder (2013), Competition between high-frequency traders and market quality, Working Paper.
13 For more information on this approach see, for example, J Brogaard, T Hendershott and R Riordan (2014), High-frequency trading and price discovery, The Review of Financial Studies 27 (8), pp 2267-2306; or J Brogaard (2011), High-frequency trading and volatility, Working Paper.
14 Market orders are orders where the given number of securities should be traded immediately at the currently most favourable rate or as soon as possible. A market order expresses the contractor's preference for time over price. Limited orders are orders to buy (sell) a security at a given price or below (above) it. The purpose of these conditional orders is to give traders protection against a transaction at unfavourable conditions. A limit order thus expresses the contracting party's preference for price over time.

In times of heightened market stress, HFT participants increasingly pursue momentum strategies …

The next step will be to study which market players tend to, on average, go in line with market movements (momentum strategy) or contrary to the market (contrarian strategy). Active traders mainly follow momentum strategies in the Bund future market, whereas passive traders, on average, buck the price trend. Consequently, passive orders dampen volatility – irrespective of the current market setting: in both the heightened-volatility and low-volatility weeks, active traders follow a momentum trading strategy, while passive market makers show contrarian trading behaviour. However, HFT participants change from contrarian strategies in calm market phases to momentum trading strategies during periods of heightened market stress.

… and, in addition ramp up their trading activity with increased volatility

A further analysis of the two weeks shows that both active and passive HFT participants ramp up their trading activity in the high-volativity week, whereas in the calm week their trading activity remains relatively constant across various levels of volatility.[15] Slower traders (non-HFT participants) do just the opposite: in the high-volatility week, they reduce their trading in inverse proportion to the increase in volatility. One reason for the positive relationship between volatility and active HFT participants' trading activity could be that these traders see more opportunities for short-term gains in times of major price swings. Another possible explanation is that HFT participants acting in the direction of the market are the cause of the higher volatility. In the absence of a specific exogenous event or an appropriate statistical instrument,[16] it is impossible to identify an exact causality.[17] In the case of passive HFT participants, the direction of causality between their trading activity and the volatility is likewise unclear. It is important to note that these traders are more likely to exercise their limit orders in periods of volatility. It is accordingly unclear whether passive HFT participants prefer an environment of heightened price volatility or whether the heightened volatility simply leads to more orders being placed, particularly if they are close to the best bid-ask prices. Owing to this endogeneity problem, at this juncture it is possible only to establish a positive relationship between the activity of active and passive HFT participants and volatility, but not a specific causal relationship.

Causal link between HFT and volatility not definitively established, however

In conclusion, active HFT participants who trade in the direction of the price movement predominate in the more turbulent market phase and increase their activity as volatility rises. This results in an increased risk of HFT participants even amplifying excessive price volatility in times of jittery markets.

Example of a directional strategy – news trading

Speed is very advantageous when responding to news

A subject of the study was "news trading", which means trading in response to important news. Since this is a direct response to the announcement of news, news trading is of specific relevance to active HFT participants, whose orders are executed immediately. Since HFT participants can make particularly good use of their speed advantage, news trading strategies are widely used among them.

Announcement of macroeconomic data particularly relevant in the case of government bonds

It is known that the publication of important macroeconomic data plays a major role in the markets for government bonds. Yields in European bond markets often respond the strongest just on either side of the announcement of

15 This result is largely confirmed by empirical studies in other markets (see E Boehmer, K Fong and J Wu (2012), International evidence on algorithmic trading, Working Paper; A Chaboud, B Chiquoine, E Hjalmarsson and C Vega (2014), Rise of the machines: Algorithmic trading in the foreign exchange market, Journal of Finance, 69 (5), pp 2045-2084). The realised variance, which represents the sum of the squared yields per time interval (here: 1 minute), is used here as a measure of volatility.

16 The objective of a statistical instrument is to have maximum correlation with the explanatory variable but to be uncorrelated with the dependent variable in order to avoid endogeneity problems. In order to measure, for instance, the impact of HFT on volatility, the instrument should be highly correlated with HFT activity but be as uncorrelated as possible with market volatility.

17 See, for example, E Benos and S Sagade (2012), High-frequency trading behaviour and its impact on market quality: evidence from the UK equity market, Bank of England Working Paper; or J Brogaard (2011), op cit.

Eurex datasets

The empirical studies outlined are based on two different datasets for DAX and Bund futures. The first dataset comprises a week of relatively high volatility, 6-13 March 2014 (during the Crimean crisis), and a week of relatively low volatility, 3-10 June 2014, when the DAX reached a new all-time high of over 10,000 points. These two volatility phases were determined using the DAX's implicit volatility index (VDAX). In the March week during the Crimean crisis, the VDAX reached a level approaching 20, with a mean of 17.2, compared with an average level of 14.1 in the preceding six months. In the June week, the VDAX was at an average level of 13.4, compared with a mean value of 16.7 over six months. The second dataset comprises individual trading days between July 2013 and June 2014 on which significant macroeconomic news was published. These are the publication dates for the ECB Governing Council's interest rate decisions, ie the first Thursday of each month in the year under review, and for US labour market data, ie the first Friday of each month. The latter refers to nonfarm payroll data, meaning those newly created jobs in the US economy each month which have the greatest effect on yields in the global sovereign bond markets. The datasets were supplemented by individual, highly volatile trading days. Each individual trading day includes all order book activities, ie all the relevant information on conducted transactions as well as bid-ask prices and volumes, plus modifications and order cancellations. The timestamp of the transactions and other order book activities is given in microseconds; all activities are additionally assigned a ranking position within these microseconds. Since existing empirical studies, on the other hand, are usually based on milliseconds or a lower frequency (eg Brogaard *et al* (2013), Gao and Mizrach (2013)), this extremely high data frequency enables particularly granular analysis.

Another key feature of the Eurex data is that they contain an HFT (high-frequency trading) identifier, which has the value 1 for HFT traders and 0 for non-HFT (NHFT) traders. It is based on a method developed by Eurex and used exclusively within the organisation. The idea of this is to investigate which traders exhibit HFT-typical behaviour on the basis of incoming orders placed by individual trading participants. First, a theoretical distribution of incoming transactions is determined on the basis of all the observed transactions; this shows, for example, the average number of orders entering the trading system per day at particular intervals from each other. Finally, this is compared with the actual distribution of transactions by each trading participant. If the expected values are significantly exceeded within the smallest time intervals (at the microsecond level), the participant is classified as an HFT trader. Therefore, the Eurex method's classification principle is based on the observed frequency of consecutive transactions by individual trading participants. The participants are only classified as HFT traders if these values are considerably higher than the average expected value overall. The group of NHFT traders therefore encompasses traditional investors as well as those computer algorithms which exhibit longer time intervals between orders submitted than HFT traders do. The resulting individual allocation of an HFT identifier may deviate from the HFT identifier based on legal requirements.

the monthly US labour market data and ECB Governing Council interest rate decisions. The behaviour of HFT participants upon the publication of these monthly Bund future market data between July 2013 and June 2014 will therefore be examined in greater depth below.[18] An additional advantage of this analysis compared with previous studies of HFT strategies in various market phases is that the announcement of macroeconomic news can be regarded as an exogenous event[19] and is generally a source of high short-term volatility. In addition, there is little empirical literature on the behaviour of HFT participants on either side of the announcement of data.[20]

Bund future yield responds strongly to US labour market data

First, HFT activity seconds before and after publication of the US labour market data is examined relative to average activity on each respective date.[21] A more detailed statistical analysis encompasses both US labour market data and the ECB Governing Council's interest rate decisions (see the box on pages 47 to 49). The chart on page 44 shows that post-announcement Bund future volatility rises sharply, on average, but already reverts to normal after just a few seconds. A clear difference between HFT and non-HFT participants in terms of behaviour emerges. Liquidity-consuming HFT participants are between around 10% and 30% more active just before and after the announcement of data than the daily average. By contrast, liquidity-consuming non-HFT participants increasingly pull out of the market, possibly because they are aware of their slower speed in response to the news. While passive HFT participants become more active only after the announcement, passive non-HFT participants likewise display above-average activity around the publication of data. However, this is due to the fact that their limit orders can adapt quickly enough to the information and are ultimately executed by the more aggressive orders issued by active HFT participants.

The next step will be to examine whether HFT participants trade in the direction of, or contrary to, market movements after the announcement of news. Surprisingly high employment gains in the United States tend to trigger expectations of rising inflation and policy rates. Therefore, in keeping with the international comovement of interest rates, Bund future prices will generally also fall in response to good employment figures and rise whenever these figures deviate negatively from the expected value.[22] The chart on page 45 shows the order flow (net order volume, ie the number of buy orders minus the number of sell orders) of HFT and non-HFT participants as well as the yield on Bund futures just before and after surprisingly positive US labour market data.[23] According to this chart, active HFT participants trade in the direction of the market in the second of the announcement and show above-average trading activity. Therefore, it stands to reason that their activity is part of the cause of the strong initial price movement. As sales and purchases by passive HFT participants are balanced in this period (order flow close to 0), HFT participants are net liquidity consumers, rather than providers, in the volatile phase owing to the publication of data.[24] The chart on page 45 illustrates that passive non-HFT

Majority of HFT participants move in direction of markets after the data are announced, …

… and thus HFT participants in this setting are net consumers of liquidity …

18 During this period, US labour market data were announced on the first Friday of each month and the decisions of the ECB Governing Council on the first Thursday of each month at a previously known time.
19 See A Chaboud, B Chiquoine, E Hjalmarsson and C Vega (2014), op cit.
20 Only one comparable study exists for the US Treasury markets: G Jiang, I Lo and G Valente (2013), High-frequency trading around macroeconomic news announcements: Evidence from the US Treasury market, Working Paper. Further macroeconomic data studies for other markets include A Chaboud, B Chiquoine, E Hjalmarsson and C Vega (2014), op cit, in the foreign exchange market and J Brogaard, T Hendershott and R Riordan (2014), op cit for the US equity market.
21 The graphic analysis incorporates only labour market data since breaking down macroeconomic data into positive and negative surprises is more clear-cut than breaking down interest rate decisions along the same lines.
22 In the underlying study, the median of analyst estimates reported by Bloomberg is used as the expected value.
23 The result is less meaningful for surprise bad news, as a result of, inter alia, outliers and a small number of observations.
24 This result runs counter to the study by J Brogaard, T Hendershott and R Riordan (2014), op cit, for the US equity market, according to which, in the aggregate, HFT traders dampen volatility owing to the preponderance of passive market makers.

Trading activity upon release of US labour market data*

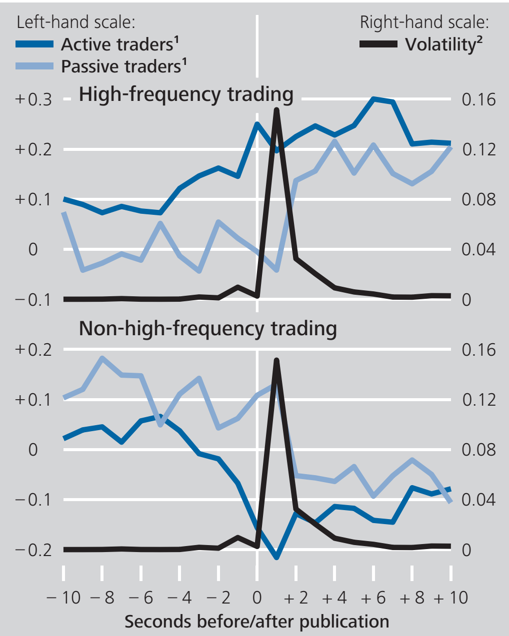

Sources: Eurex, Bloomberg, and Bundesbank calculations. * Monthly publication of nonfarm payrolls. **1** Share of the trading volume in relation to the total trading volume per second. The surplus relative to the entire trading day is shown here. **2** Surplus relative to the average of the entire trading day in %.
Deutsche Bundesbank

participants occupy the other side of the positions of active HFT participants.

… and can generate profits nearly instantaneously

The statistical analysis (see the box on pages 47 to 49) supports the graphic results: active HFT participants respond at the second of the publication in the direction of the surprise component of the news. Following this initial response, the sign of the order flow reverses itself after a few seconds, which would suggest that the active HFT participants have already realised their trading gains.

The results show news trading behaviour which is typical of HFT, and which is based on fast and relatively aggressive trading: HFT players use their speed advantage to buy or sell Bund future contracts in response to the data at a favourable price as the new information has not yet been incorporated into the price. Once slower market participants have also responded to the news with a time lag and have amplified the initial price movements with their trades, active HFT participants then close their positions. They generate, in a very short period of time, significant gains, in some cases, depending on the strength of the market reaction.

Orders of passive non-HFT participants executed by fast HFT orders

Passive non-HFT participants issuing orders against the market trend following the announcement of news cannot adapt their orders quickly enough and thus fall victim to “adverse selection”: their orders are executed by orders of active HFT participants at prices that are not favourable to them. It can be empirically confirmed that passive non-HFT participants increasingly withdraw even before the news is announced in such market phases.

Rapid action by HFT participants leads to higher price efficiency, but also to increased short-term volatility

The trading behaviour of HFT participants is statistically significant only in the first second after the data are announced. For more detailed results regarding market players’ initial reactions, the box on pages 47 to 49 studies the publication of macroeconomic news on the basis of transactions (ticks). The results show that, on the one hand, active HFT participants contribute significantly more to price efficiency than active non-HFT participants: their activity causes new information to be incorporated into prices more quickly. On the other hand, their immediate and aggressive trading in response to price-relevant news triggers a very high short-term volatility, an “overshooting” of sorts. It should be borne in mind that prices are not made more informative by the activity of HFT participants as such. The latter would be the case if their trading were to create new information which would not be incorporated into prices without their actions.[25] Rather, the contribution made by these traders to higher price efficiency is that new information is factored into prices microseconds faster than would have been the case in their absence. However, the economic value of prices which

25 See also J Hirshleifer (1971), The private and social value of information and the reward to inventive activity, The American Economic Review, 61, pp 561-574; and T Foucault (2016), Where are the risks in high frequency trading?, in Banque de France, Financial Stability Review, No 20, April 2016.

are more efficient by fractions of a second is difficult for the human observer to comprehend.

Example of a passive strategy – market making

Liquidity provision is of great importance, particularly in times of stress

A fear among regulators and market participants is that passive HFT participants provide liquidity only in calm market phases and tend to withdraw it in times of stress, where liquidity is particularly needed.[26] Such behaviour also has implications for financial stability, as in the case of market shocks a reduced supply of liquidity could even intensify the shock.[27]

When analysing the provision of liquidity in Bund futures, a distinction is made between a period of unexpected volatility triggered by a surprise increase in traders' risk aversion, such as the Crimean crisis in March 2014, and a phase of expected volatility, such as when the US labour market data are announced.[28] In order to assess whether the liquidity increases or decreases at any given point in time, a deletion ratio, defined as the ratio of orders deleted from the order book to new orders over a given time period.[29]

HFT participants increasingly delete orders in periods of more highly jittery markets …

HFT participants exhibited a deletion ratio of 77% in the volatile week in March 2014, which was higher than in the calm week in June 2014, at 72%; for non-HFT participants, by contrast, there is virtually no discernible difference. This suggests that the provision of liquidity by non-HFT participants is more constant across various market phases.

… and provide less liquidity as volatility increases

Furthermore, one sees that a surprise increase in volatility in the turbulent week in March led HFT market makers to increasingly cancel their posted limit orders.[30] Non-HFT participants did the opposite: when volatility was higher, they withdrew fewer orders. In the calmer trading week in June, both HFT and non-HFT participants tended to delete fewer orders as volatility increased. As the week in June was a period in

Trading behaviour upon release of US labour market data*

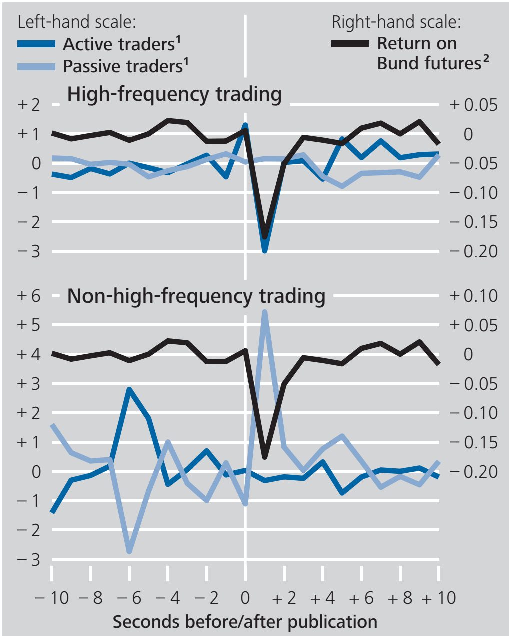

Sources: Eurex, Bloomberg, and Bundesbank calculations. * Monthly publication of nonfarm payrolls. **1** Net trading volume per interval of one second. The surplus relative to the average of the entire trading day is shown here. **2** Surplus relative to the average of the entire trading day in %.
Deutsche Bundesbank

which markets were not very jittery, a slight uptick in volatility during such a period does not appear to be cause for HFT traders to withdraw

26 See European Securities and Markets Authorities (ESMA) (2011), Report of trends, risks, and vulnerabilities, European Securities and Markets Authorities, Working Paper; and Australian Securities and Investments Commission (ASIC) (2012), Report 331: Dark liquidity and high-frequency trading, Working Paper.
27 See Y Ait-Sahalia and M Saglam (2014), High-frequency traders: taking advantage of speed, Working Paper.
28 See K Schlepper (2016), op cit.
29 On the basis of order deletions, a more detailed analysis of the impact of volatility on the provision of liquidity by market players can be conducted than is possible based on transaction data. For transaction data (see eg A Chaboud, B Chiquoine, E Hjalmarsson and C Vega (2014), op cit, and J Brogaard, T Hendershott and R Riordan (2014), op cit) the traded volume is used to measure the provision of liquidity. However, as this is highly correlated with volatility, previous studies often suffer from an endogeneity problem. For example, E Benos and S Sagade (2012), op cit, argue that, in periods of higher volatility, passive limit orders are more likely to be carried out, thus automatically increasing the trading volume.
30 To this end, a dummy variable is created which has a value of 1 if the deletion ratio exceeds the average figure for the whole week during the interval (1 minute) and 0 otherwise. This variable was regressed on the volatility in the preceding minute using a probit method.

Liquidity provision around the time of publication of US labour market data

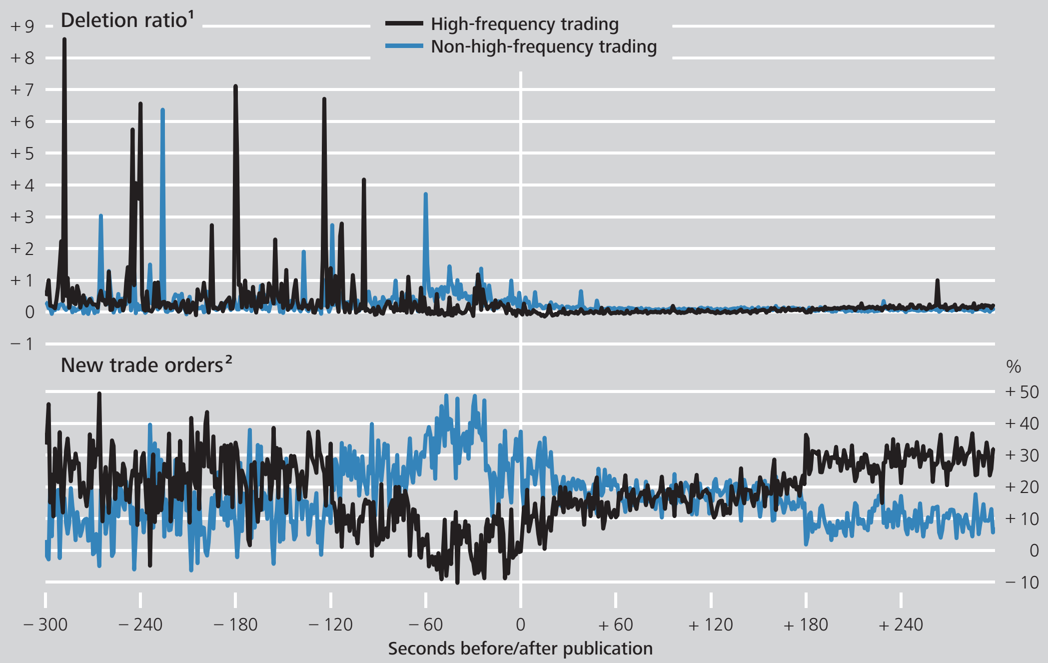

Sources: Eurex, Bloomberg, and Bundesbank calculations. **1** Ratio of orders deleted from the order book to new orders per interval of one second. The surplus rate produced by the difference from the average of the entire trading day is shown. Values greater than 1 mean that more orders were deleted in the given second than new orders placed. **2** Ratio of new trade orders relative to all order book activities (ie new orders, transactions, deletions and modifications). The surplus rate produced by the difference from the average of the entire trading day is shown.

Deutsche Bundesbank

their liquidity from the market. The fact that they responded much more sensitively to market fluctuations in the more turbulent week could indicate the existence of certain volatility thresholds which, if overshot, cause HFT traders to deem the market too risky and to increasingly pull out. In addition, during such periods HFT traders have no information advantage over non-HFT participants, unlike, for example, is the case regarding the announcement of news, which HFT traders can process faster than slower market participants. Thus, from a certain stress level, HFT participants increasingly cut back their supply of liquidity in order to avoid higher hedging costs when performing their market-making activities.

HFT traders also delete an excessive number of trading orders before the announcement of US labour market data

Furthermore, the empirical results show that both HFT and non-HFT participants delete an excessive number of orders compared to the average on that trading day even minutes before the publication of US labour market data – and thus prior to the period of expected volatility. For HFT traders, however, the deletion ratio is significantly higher than that for slower traders. Looking at the new orders relative to all order book activity on either side of the announcement of US labour market data, the order rates of HFT and non-HFT participants can be seen to diverge (see the chart above): while HFT traders place an above-average number of orders in the minutes prior to the announcement, activity decreases just prior to the announcement. At the same time, orders from non-HFT participants go up sharply.

The results show that, irrespective of the nature of volatility – expected or unexpected – HFT participants provide less liquidity as market stress rises and increasingly pull out of the Bund future market. As the increase in the volatility is anticipated *ex ante* when significant macroeconomic news is published, the orders are already increasingly deleted prior to the announcement, whereas if the rise in volatility

Empirical evidence of HFT participants' reaction to the publication of important news

Especially when important news such as US labour market data and ECB Governing Council interest rate decisions is published, high-frequency trading (HFT) participants can use their speed advantage, and react within the first few seconds after the announcement. On some days when US labour market data are published, up to around 500 transactions may take place during the second in which the announcement is made. To obtain a more accurate picture of HFT participants' initial reaction, which generally takes place within milliseconds or even microseconds, the following statistical analysis is performed at the tick level. This also reduces potential endogeneity problems because individual transactions can be analysed sequentially.[1] The following vector autoregressive model (VAR model) with 10 lags is used to capture market participants' reactions to the announcement of US labour market figures and ECB interest rate decisions and to measure the effect on the Bund future return:

$$r_t = \alpha + \sum_{i=1}^{l} \beta_i OF_{t-i}^{hft} + \sum_{i=1}^{l} \gamma_i OF_{t-i}^{nhft} + \sum_{i=1}^{l} \delta_i r_{t-i} + \sum_{i=0}^{l} \psi_i D_i + \epsilon_{1t}$$

$$OF_t^{hft} = \kappa + \sum_{i=1}^{l} \eta_i OF_{t-i}^{hft} + \sum_{i=1}^{l} \lambda_i OF_{t-i}^{nhft} + \sum_{i=1}^{l} \nu_i r_{t-i} + \sum_{i=0}^{l} \phi_i D_i + \epsilon_{2t}$$

$$OF_t^{nhft} = \zeta + \sum_{i=1}^{l} \rho_i OF_{t-i}^{hft} + \sum_{i=1}^{l} \tau_i OF_{t-i}^{nhft} + \sum_{i=1}^{l} \nu_i r_{t-i} + \sum_{i=0}^{l} \pi_i D_i + \epsilon_{3t} \ ,$$

Here, the dummy variable D is entered as an exogenous variable; it takes the value of 1 at the time data is released and 0 otherwise. OF denotes the net order flow of HFT and non-HFT, and r is the log of the Bund future return at time t. The other variables are coefficients estimated using the maximum likelihood method. A Cholesky decomposition is applied to transform the VAR model into a vector moving average (VMA) form which can be used to identify the impact of HFT and non-HFT activity on returns:

$$\begin{pmatrix} r_t \\ OF_t^{hft} \\ OF_t^{nhft} \end{pmatrix} = \begin{pmatrix} a(L)b(L)c(L) \\ d(L)e(L)f(L) \\ g(L)h(L)k(L) \end{pmatrix} \begin{pmatrix} e_{1,t} \\ e_{2,t} \\ e_{3,t} \end{pmatrix} + \begin{pmatrix} q(L) \\ r(L) \\ u(L) \end{pmatrix} \begin{pmatrix} X_{1,t} \\ X_{2,t} \\ X_{3,t} \end{pmatrix}$$

The standard errors are orthogonalised, such that $e_t e_t' = I$, allowing causal conclusions about shocks from individual elements of e_t. The polynomials $a(L)$ to $k(L)$ represent the impulse-response functions of the three variables to shocks, while $q(L)$ to $u(L)$ capture the cumulative impact of the dummy variable on the three dependent variables. Furthermore, $b(L)$ and $c(L)$ represent the impulse-response functions of the HFT and non-HFT order flow, respectively, and can be interpreted as permanent price effects of an innovation in the order flow of HFT and non-HFT.

The upper table on the following page shows the effect that ECB interest rate decisions and NFP releases have on the order

1 When using time intervals, a large number of transactions occur simultaneously (ie within one second, for example). The number of transactions also varies strongly, ranging from five transactions in the case of an ECB interest rate decision up to 1,330 transactions following a US labour market report in the second after publication. A tick-based analysis also allows the variability to be taken into account, which is not possible in an analysis based on fixed intervals.

The effect of news on trading behaviour

Item	Active HFT	Active NHFT	Passive HFT	Passive NHFT	HFT	NHFT
Short-term effect (tick = 0)	– 0.089	0.004	– 0.019	0.065	– 0.099	0.096
Standard error	(0.304)	(0.237)	(0.16)	(0.346)	(0.292)	(0.316)
Long-term effect (tick = 10)	– 6.715***	– 0.875	0.967	5.686***	– 5.892***	5.709***
Standard error	(1.764)	(1.228)	(0.654)	(1.89)	(1.704)	(1.7)
Long-term effect – short-term effect	– 6.626***	– 0.879	0.986	5.621***	– 5.793***	5.613***
Standard error	(1.79)	(1.251)	(0.673)	(1.921)	(1.729)	(1.729)

* Statistically significant at 10%, ** at 5% and *** at 1%.
Deutsche Bundesbank

The effect of trading behaviour on the Bund future return

Item	Active HFT	Active NHFT	Active HFT – NHFT
Short-term effect (tick = 0)	0.069**	0.081***	– 0.012
Standard error	(0.031)	(0.032)	(0.045)
Long-term effect (tick = 10)	0.607***	0.120*	0.486***
Standard error	(0.063)	(0.067)	(0.092)
Long-term effect – short-term effect	– 6.626***	– 0.879	– 5.793***
Standard error	(0.07)	(0.074)	(0.102)

* Statistically significant at 10%, ** at 5% and *** at 1%.
Deutsche Bundesbank

flow of HFT and non-HFT, given by $r(L)$ and $u(L)$ in the model. This shows that only active HFT and passive non-HFT participants display a significant reaction to the data release. Within the first 10 ticks, HFT participants trade in the direction of the market; this effect becomes stronger the more ticks are considered. Passive non-HFT participants show the opposite reaction because they take the other side of the active non-HFT participants' trades that are executed.

The results of the impulse-response functions for active market participants[2] (see the lower table on this page) show that both HFT and non-HFT order flows have a significant positive short-term effect on the return. While the effect expands for HFT participants over the subsequent 10 ticks (long-term effect), it decreases slightly in the case of non-HFT participants. The chart on page 49 shows the cumulative impulse-response functions (IRF) for 10 events into the future and their 95% confidence bands. The IRF of HFT participants after 10 ticks is around five times higher than that of non-HFT participants. This reveals that a shock in the HFT order flow has a 400% greater price effect than a comparable shock in the non-HFT order flow. This suggests that HFT orders contain more information than those of non-HFT participants. This finding is corroborated by the significant difference between the IRFs of HFT participants and non-HFT participants (Column 3 in the lower table on this page). Nevertheless, the stronger information efficiency found in the first 10 ticks applies only to a time-span of milliseconds. The economic benefit resulting from this seems doubtful. It is also unclear whether trading by active HFT par-

2 The analysis is confined to active market participants, since they initiate the trades and thereby contribute to price discovery. This is also consistent with the method of T Henderschott and R Riordan (2011), Algorithmic trading and information, working paper; and A Chaboud, B Chiquoine, E Hjalmarsson and C Vega (2014), Rise of the machines: Algorithmic trading in the foreign exchange market, Journal of Finance, 69(5),pp 2045-2084.

ticipants in response to news generates a purely informative contribution or whether this creates additional unwanted short-term volatility (white noise). Generally speaking, volatility can have both a permanent (information-based) component and a transitory component. These two components are therefore decomposed using the variance decomposition method developed by Hasbrouk (1991 and 1993). The aim of this is to gauge the impact of active HFT and non-HFT transactions on the share of permanent and transitory variance. According to Hasbrouck, the observed price can be written as

$$p_t = m_t + s_t ,$$

where m_t describes the permanent component and s_t denotes the transitory component or the error term. In this context, the error term may be interpreted as the lagged adjustment of the price to new information. On this basis, the variances of the permanent and transitory components can then be calculated as a function of the coefficients of the VMA process and the variance-covariance matrix.[3] We find that, at 9.2%, HFT participants make a significantly larger contribution to permanent variance than non-HFT participants (0.4%). However, through their more aggressive trading, HFT participants also contribute much more strongly to transitory variance, which, at 33.4% after 10 ticks, is more than ten times higher than the figure for non-HFT participants (just 2.3%). The information-to-noise ratio for HFT participants is therefore well below 1, indicating that their rapid trading generates more transitory than information-based variance. It can be concluded from this that, as a rule, active HFT participants "overshoot the mark" in response to macroeconomic news. This means that they do act in the direction of the market movement in line with the surprise component of the news, thus contributing to the price discovery process. Nevertheless, their reaction appears excessive measured in terms of the long-term implications of the news for the price.

Impulse-response functions* upon release of US labour market data

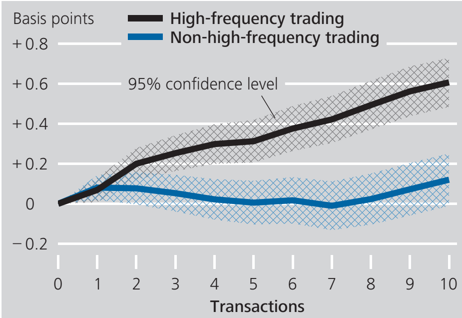

Sources: Eurex, Bloomberg, and Bundesbank calculations. * Effect of an increase by 1 standard deviation in the net volume of orders (ie the number of buy orders minus the number of sell orders) on the Bund future return.
Deutsche Bundesbank

3 For more details, see J Hasbrouck (1993), Assessing the quality of a security market: a new approach to transaction-cost measurement, Review of Financial Studies 6, pp 191-212; K Schlepper, High-frequency trading in the Bund futures market, Deutsche Bundesbank Discussion Paper No 15/2016; and E Benos and S Sagade (2012) High-frequency trading behaviour and its impact on market quality: evidence from the UK equity market, Bank of England working paper.

comes as a surprise, these deletions occur only once the markets have responded.

The importance of HFT traders in the limit order book

Limit order book as a key element for understanding the market …

One essential component for a better understanding of HFT is taking a closer look at the environment in which its HFT algorithms interact with other market players. On nearly all regular trading platforms, this interaction takes place in the limit order book (LOB). All incoming and not immediately executable limited buy and sell orders are entered into it. Market orders, ie orders which are not limited and are instead placed to be executed immediately, are not entered directly into the LOB. Market orders nevertheless have an impact on the LOB because they interact with the limited orders entered into it through orders being executed. Orders are processed in the electronic trading system of the given stock exchange according to a fixed set of rules (market model) known to the market players in terms of their priority for execution.[31]

… represents overall liquidity

At each point in time during the trading process, the LOB – the aggregate total of all unexecuted limited buy and sell orders – represents the total demand and supply sides for the financing instrument. The LOB at time t thus represents the liquidity which is available at any given time for executing transactions for market players (see chart on page 51).[32,33]

Price discovery process takes place in the LOB

A transaction takes place in the LOB if a new buy (sell) order arrives or if an existing order is modified in such a way that it overcomes the bid-ask spread and can be executed with a sell (buy) order on the other side of the LOB. Price discovery is therefore a complex process resulting from the constant individual entry, deletion and modification of orders by market participants and the ensuing transactions at the respective best price.

The processes running in the order book in the trading of a highly liquid financial instrument, such as the DAX or Bund future, are very numerous and complex and considerably increase the amount of work required for an LOB analysis.[34] From a regulatory perspective, an LOB analysis nevertheless affords a significantly better insight into the price discovery process, which is one of the most important functions of securities and futures markets and to which must be ascribed major economic relevance for the allocation of capital and for financial stability.

First-time analysis of the order book for very small time scales and taking account of HFT

A comprehensive analysis of the LOB makes it possible to investigate a number of questions concerning the German capital market about which there has so far been little economic research. An investigation is made, for example, to determine the fundamental microstructural characteristics of the LOB (such as execution times of orders) for ultra-short time scales and with special consideration of HFT orders. It is now also possible to answer the question as to where in the order book HFT traders place their orders and how large their share of the liquidity provided in the LOB is. A similarly important aspect is the dynamics of the provision of liquidity by HFT and NHFT participants over time. The LOB analysis also supplies answers to the question of whether there are particular periods in which one or even

31 In most market models, the order of priority is set initially by the price at which the investor is prepared to buy or sell the security in question. At the second level, if there is more than one order at a given price, there is a further prioritisation according to the time at which the order entered into the LOB (with older orders that arrived earlier being given priority over more recent ones).
32 Most stock exchanges grant their market players a partial or even complete insight in real time into the current status of the LOB (open order book).
33 See M Gould et al (2013), Limit Order Books, Journal of Quantitative Finance, 13 (11), pp 1709-1742; and M Paddrik et al (2014), Effects of Limit Order Book Information Levels on Market Stability Metrics, OFR working paper.
34 For the sample of DAX future data examined in the LOB analysis, the volume of a normal daily rate of data fluctuates between about 0.9 million and 7 million individual events in the order book, of which each is registered with several dozen information units (such as the limit price, the precise time, the number of contracts, the identifier as an HFT order, etc). HFT data and the related research therefore have a marked "big data" character.

Representation of an order book

Depth of the order book (number of securities offered or in demand)

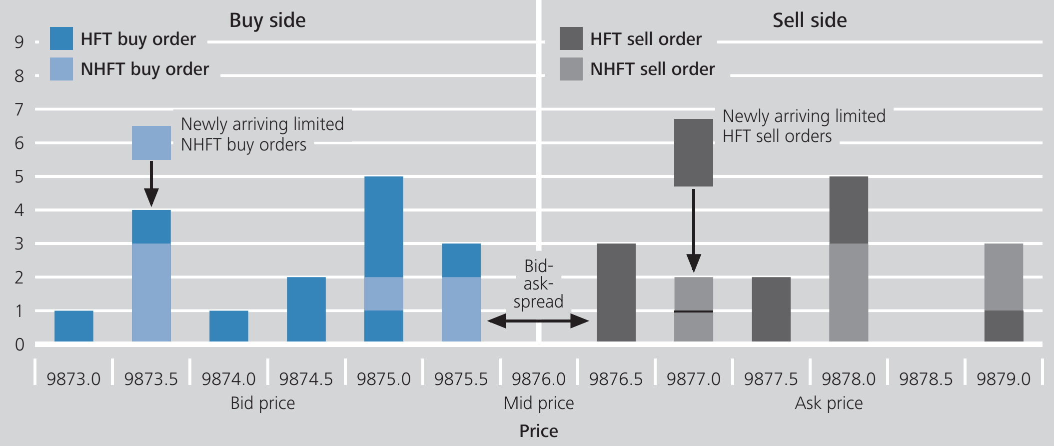

Deutsche Bundesbank

both parties significantly reduce their presence in the LOB (liquidity holes).

Both for the composition of the LOB at a given point in time and for the analysis of the LOB dynamics over time, it is sufficient if the sequence of the individual order entries, modifications and deletions as well as partial or full executions are processed in a way that is consistent with the set of rules of the trading platform. Using the available Eurex order data and the HFT identifier, it is therefore possible to reconstruct the situation in the LOB at virtually any time down to the level of a microsecond.

Selection of a small but meaningful dataset

From the available Eurex data sample, a small selection of 12 individual trading days from 2013 to 2015 is made for the DAX future. The selection of trading days is guided by the requirement that the sample should include both normal days characterised by calmer trading and those marked by higher intraday volatility and dynamic, strongly news-driven market activity.[35]

Characteristics of the LOB

Just under half of all activities in the order book are HFT

The 12 selected trading days of the DAX future contain a total of around 21.1 million order activities. Of the roughly 1.75 million daily LOB events, 52.2% are due to HFT traders and 47.8% to NHFT participants. For the actual contract volume, the figures are 41.3% for HFT and 58.7% for NHFT. On average, an executed HFT order is 1.31 contracts in size. For an NHFT order, the figure is 1.68.

Analyses for very small time scales call for special method

The image of steady, continuous market activity on time scales such as minutes or seconds becomes more and more discrete and unsteady when progressively “zooming into” shorter time scales in the sub-second range.[36] The granularity of market activity at the millisecond or microsecond level therefore plays a big part in the analysis of high-frequency order book data. A continuous time flow with an approximately steady level of activity on time scales that can be perceived by human beings is thus often transformed into a discrete sequence of

35 The days studied in the LOB analysis are 5 July 2013, 2 August 2013, 6 September 2013, 2 October 2013, 8 November 2013, 6 December 2013 (NFP days), 3 to 6 June 2014, 9 June 2014 (normal days), and 7 May 2015 (very high volatility). For four of these days (2 August 2013, 6 September 2013, 3 and 10 June 2014), the LOB is reconstructed in an extensive analysis down to the level of a microsecond.
36 It is thus not unusual, for example, for there to be periods at the level of one second in which several hundred orders or price discoveries occur within the space of a second. Conversely, there are many seconds in which no activities whatsoever occur in the LOB. A one-second volatility has widely differing implications in such circumstances.

Average time span before a new LOB activity*

in seconds

LOB activity	HFT	NHFT
New order	0.0143	0.0436
Modification	0.0132	0.0256
Deletion	0.0118	0.0189

* Times between a change in the bid-ask price or in the contract quantities offered or in demand at this price and a new LOB activity.

Deutsche Bundesbank

"activity clusters" in the sub-second range. In order to take account of this characteristic, all calculations of parameters, such as volatility or returns, are calculated at the level of the individual ticks in the order book on the basis of an "event to event" procedure.[37]

The response times to changes in the LOB differ significantly between HFT and NHFT traders. After a change in the bid-ask price or in the contract quantities offered or in demand at these prices, HFT participants respond significantly more quickly than NHFT traders, irrespective of whether the activity initiated next in the LOB is a new order, a modification or a deletion.

Time until an HFT order is executed very much shorter than in the case of NHFT

Major differences are likewise apparent in terms of the amount of time a limit order that is not executed immediately spends in the LOB until it is executed. The median time until half of all HFT orders are executed amounts to 1.15 seconds. For NHFT, the median lies at 6.02 seconds.[38]

HFT liquidity near to the best bid-ask price

The assertion that passive HFT traders make an essential contribution to market liquidity is a key argument of the proponents of HFT. Owing to their presence, so the reasoning goes, other market players can trust in their orders being executed almost permanently at fair prices and within a short period of time. This argument thus implies that HFT in the vicinity of the respective best bid-ask price (eg up to 3 ticks from the current DAX future price) provides significant liquidity so that even a fairly large order of an NHFT trader can be executed without an overly large adverse price movement for the trader.[39] Furthermore, this adequate liquidity should be largely permanently present. With the LOB analysis, this assumption, in addition to the results from studying the Bund future data, is to be examined directly using a further method.

Its proponents see HFT as a major guarantor of liquid trading

To this end, it is calculated how many DAX contracts in the LOB are being offered or are in demand at the individual prices at each point in time by both HFT and NHFT traders. Additionally, a focused analysis of liquidity provision is conducted near to the best bid-ask prices in

37 With regard to the sampling frequency of major parameters, such as returns or volatility, a distinction can be made between two basic procedures. In the case of the former, use is made of calculation periods that are spread evenly over time, where there is, for example, always one second between consecutive sampling points. These produce time series with a one-second time scale. This procedure is chosen here mainly in the Bund future analyses. In the case of calculation periods based on the actual arrival times of orders in the LOB, the time series are based on an "event to event" sampling. The time intervals between the individual observation points are mostly variable and are produced by the (stochastic) arrival patterns of orders in the LOB. Sampling of this kind is often more suitable for taking account of the discrete granularity of the market in the sub-second range. In a variant of this procedure, parameter measurements are always carried out at the time of actual transactions. This is referred to as a time series sampled "trade to trade". Both variants are applied in the LOB analysis of the DAX future data.
38 The averages are 61 seconds for HFT and 369 seconds for NHFT. The major difference between the average and the median is that the distribution of the execution times follows an exponential function. In this case, the median is therefore to be seen as the more meaningful figure.
39 The minimum tick size is 1 tick, ie the smallest unit at which the asset is priced on the trading platform. In the case of the DAX future, 1 tick amounts to 0.5 DAX point. In other words, if, say, the best bid-ask price is currently 9876.5 points as in the chart on p 51, the 3-tick group for HFT traders on the sell side comprises all contracts offered by them between 9876.5 and 9877.5 points (seven contracts in the example).

Average HFT liquidity provision in the LOB

Item	Total LOB HFT	20 ticks HFT	10 ticks HFT	5 ticks HFT	3 ticks HFT
	%				
2 August 2013	55.9	35.4	32.9	42.2	47.6
6 September 2013	70.2	38.9	37.0	44.4	49.0
3 June 2014	57.6	35.7	34.4	40.9	42.8
10 June 2014	52.5	34.5	35.6	40.8	44.0
Average	59.0	36.1	34.9	42.1	45.8
	in € million				
Average HFT share	603.3	46.9	22.2	11.6	6.7

Deutsche Bundesbank

each case, ie where price discovery and trading activity effectively take place. Besides the absolute number of contracts, we thus also study how the liquidity of the two types of traders in the LOB is distributed with increasing distance from the bid-ask price. To do this, the aggregated contract volumes of the orders standing in the LOB within 3, 5, 10 and 20 ticks from the best bid-ask price are determined in each case for the buy and sell sides and the percentage due to HFT is calculated. This percentage represents the HFT liquidity in these LOB buckets and allows a more precise picture of the average provision of liquidity over the course of the day. Building on this, the temporal dynamics of the liquidity provision are analysed.[40]

HFT provides liquidity in all areas of the order book …

The LOB analysis confirms that HFT contributes a significant percentage of the liquidity provided. The HFT-induced liquidity is distributed broadly over the entire LOB and is by no means concentrated only on the area in the immediate vicinity of the best bid-ask price (see the adjacent chart).

With regard to the distribution of HFT liquidity, it becomes clear from the table above that the

Share of HFT liquidity provided in the DAX future order book

Percentage share of liquidity

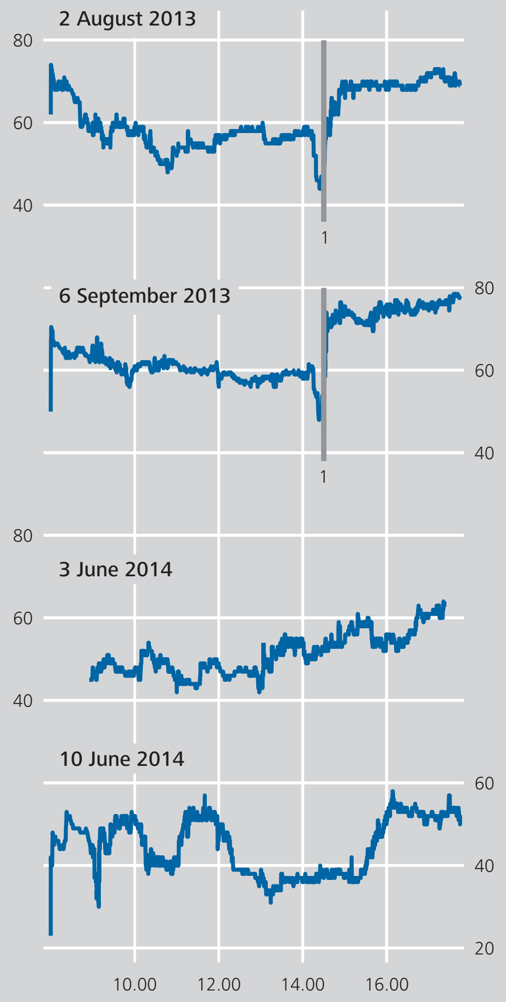

Sources: Eurex and Bundesbank calculations. **1** Publication of US labour market data.
Deutsche Bundesbank

40 The trading days selected for this more in-depth analysis are 2 August 2013, 6 September 2013, 3 and 10 June 2014.

Share of liquidity provided by HFT traders near to the best bid-ask price*

Percentage share of liquidity

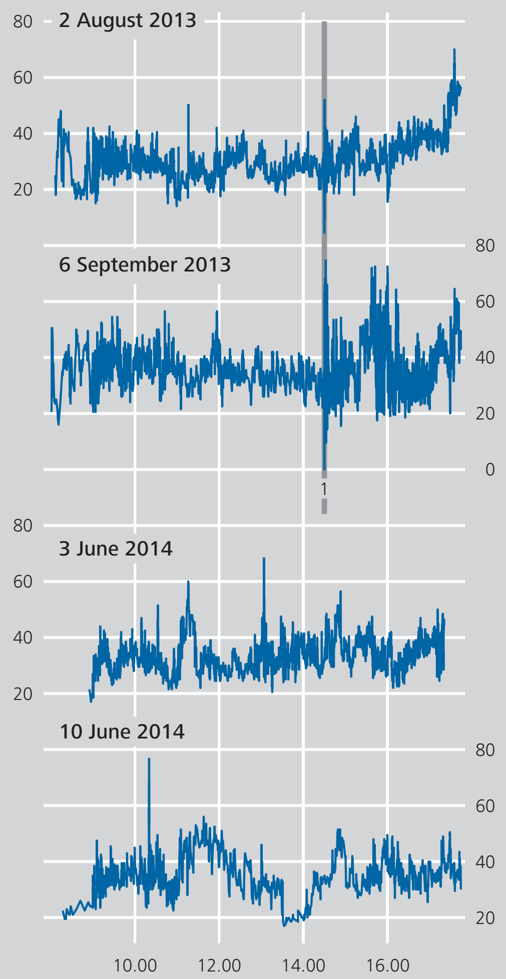

Sources: Eurex and Bundesbank calculations. * Up to 5 DAX points away from the best bid-ask price. **1** Publication of US labour market data.

Deutsche Bundesbank

… but does not do so everywhere to the same extent …

HFT percentage in the immediate vicinity of the best bid-ask price (3 and 5 ticks) is quite high at around 45% and then decreases markedly in medium distance (10 to 20 ticks) to the best bid-ask price. In this medium range of the LOB, the contribution of HFT, for instance, at 10 ticks (=5 DAX points) distance is significantly lower at no more than just under 35% (see the chart above). Further away from the best bid-ask price, the HFT liquidity then increases quite sharply, making the average figure for the LOB as a whole rise to almost 60%.

The existence of high HFT liquidity far away from the best bid-ask price appears surprising at first sight, as limit orders placed in this range stay for a considerably longer period of time until they are executed and the HFT advantages of speed do not come directly into play in the case of these orders. It would therefore be expected that HFT orders play a particularly active role, above all, in the immediate vicinity of the best bid-ask price.

… and also not always only as the market maker

One explanation might be that the HFT liquidity far from the best bid-ask price consists of orders by HFT traders that pursue passive strategies which are not, however, based on any market-making concept.[41] The further rise in the HFT percentage in the vicinity of the best bid-ask price is, by contrast, consistent with the expected behaviour of HFT and can be explained by market-making strategies.

HFT market makers confirmed as major providers of liquidity

The assertion that HFT provides other market players almost permanently with important liquidity close to the best bid-ask price can therefore largely be confirmed. In normal times, HFT therefore supports liquidity.

HFT liquidity is not granted all the time, however …

The timing of liquidity provision by HFT traders proves to be quite stable for much of the times on the days under consideration. As was already becoming apparent in the study of the Bund future data, the liquidity situation also changes quite clearly for a time in the LOB analysis of the DAX future, however, if there is a strong, anticipated market event such as the NFP announcement (see the chart on page 59). It is true that a reduction in liquidity can also be noted in the case of NHFT traders, but its scale is not as great as in the case of HFT.

41 Orders entered early take a privileged place in the sequence of priority at this limit price. This means that the investors secure themselves a "good seat" in the LOB at this price. If the best bid-ask price should then move near to the entered order at a subsequent point in time, HFT traders can use their speed advantage in analysing the market situation and the execution of orders in order to wait until the last moment to decide whether they want to maintain (and perhaps execute) the order or whether they simply wish to delete it shortly beforehand. Such behaviour might also explain part of the high number of deletions of HFT orders.

Unusual activity patterns among HFT participants

In the debate surrounding high-frequency trading (HFT) participants, one concern that has repeatedly been raised is that a number of HFT participants might be taking advantage of their superior speed to run trading strategies that are unfair on other market participants and that might not constitute market-compliant behaviour.[1,2] However, in the absence of robust empirical data, this discussion has so far largely been driven by what is believed to be anecdotal evidence. One such piece of evidence that is cited time and again is the high incidence of order cancellations. Eurex trading data can be used to explore whether order cancellations exhibit unusual activity patterns that might point to incorrectly programmed algorithms or to behaviour on the part of individual participants that is not market-compliant. The results indicate that cancellations do indeed exhibit some irregularities for which there is no straightforward explanation.

In the DAX Futures dataset covering 12 days, order cancellations accounted for a total of 5.7 million of the roughly 21.1 million order activities.[3] This shows that no unusual behavioural patterns are immediately obvious for the vast majority of the cancellations. Only a small quantity (belonging to a previously unknown category) of cancellation activities are striking. Each of them is a rapid and repeated sequence made up of the entry and almost immediate cancellation of limited low-volume buy (or sell) orders at the current best bid (or ask) price.[4] This pattern of entering and immediately cancelling orders in the limit order book (LOB) takes place as a rapid repetitive cycle, with most orders being cancelled in less than a millisecond, only to then be re-entered identically in the LOB almost as quickly.[5] This cycle is then repeated up to several dozen times without any changes being made to order parameters such as the limit price or contract quantity. These events will be referred to in the remainder of this box as "rapid order entry and cancellation cycles" (ROECCs) (see the chart on page 56).

In the DAX Futures sample, ROECC events that have three or more directly related cycles and a mean lifetime of their individual orders of 1 millisecond or less, are by no means a rare phenomenon, with 4,882 observations, and occur several hundred times a day in the DAX Futures market.[6]

Around 96.5% thereof – 4,711 in number – were generated by HFT participants.[7] In

1 The HFT debate gained particular prominence when Michael Lewis's book "Flash Boys" was published in 2013. It is worth noting here that the criticism which that book directs at HFT participants, even if it were true, would primarily only be applicable to the United States and the exchange system in operation there. Given that Germany's capital market has a different exchange landscape altogether and a procedure of its own for transmitting orders, the systematic front running described so sensationally in "Flash Boys" would hardly be possible in Germany.
2 Market-compliant behaviour, for the purposes of this box, is any trading activity that conforms to the exchange rules and does not violate applicable statutory provisions, notably the EU's Market Abuse Regulation (Regulation (EU) No 596/2014, or MAR) and Market Abuse Directive (2014/57/EU, or CRIM MAR).
3 A median of 1.22 seconds lapses between order entry and cancellation. The average is 123 seconds. Here, too, the exponential distribution is striking.
4 In around 94% of cases, the order size (just one contract) is the minimum order size in the DAX Futures market. For the span of its lifetime in the LOB, that order represents the current best bid or ask price.
5 Fairly typical ROECC times are order lifetimes of 30 to 50 microseconds, say, followed by a break of approximately 280 microseconds until the order is re-entered. The fastest cancellations come after around 11 microseconds.
6 For the purposes of this analysis, the minimum number of cycles from which an event is counted as an ROECC was set at three. The subsequent analyses were repeated with values of five and eight, however, without producing any systematic changes in the results.
7 The mean number of cycles is 5.7 repetitions, and the most frequently observed ROECC event of a single HFT participant was made up of 91 cycles. In median terms, an ROECC order is cancelled after approximately 220 microseconds and re-entered roughly 4 milliseconds later. Thus, ROECC events rarely last longer than approximately 25 milliseconds, with many also taking place in the sub-millisecond range.

Example of rapid order entry and cancellation

Number of rapid order entry and cancellation cycles

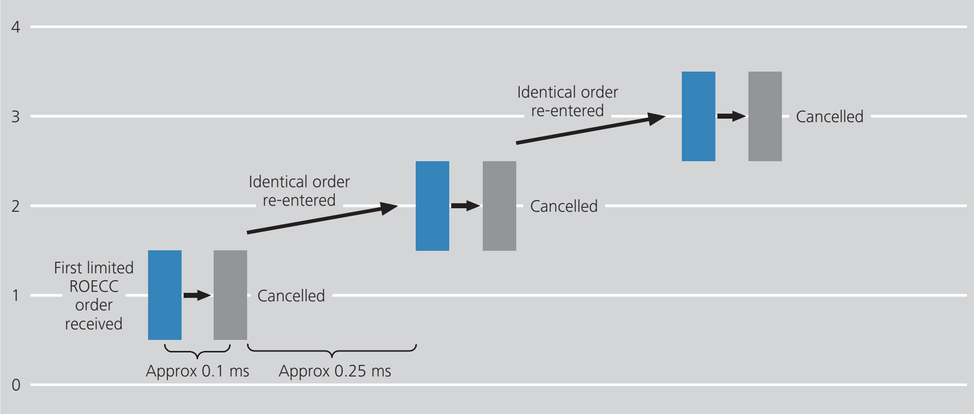

Deutsche Bundesbank

almost all cases, there is no other activity in the LOB during this time. Likewise, order executions during the course of an ROECC event are very rare.

There is no straightforward explanation for the behaviour pattern revealed by these ROECC events. The HFT algorithms responsible for them, however, clearly respond to the bid-ask spread in the LOB. That makes sense if one assumes that an ROECC order is not designed to be executed in the first place (see the chart on page 57).[8]

A first possible explanation for the ROECC events identified here could be that algorithm coding errors or self-trading are to blame.[9] Errors are a somewhat unlikely cause, however, given the evidence of a systematic response to changes in the bid-ask spread as well as the repeated incidence over a period of at least a year. Self-trading, on the other hand, is a plausible explanation that would be compatible with market-compliant behaviour since it does not constitute deliberate misconduct.

Another possible explanation is the "tracking hypothesis", according to which liquidity providers cancel and replace their orders in quick succession in an effort to keep pace with rapidly changing LOB conditions ("price chasing"). This explanation can be ruled out in the ROECC cases observed here because each subsequent order is absolutely identical with the previous ones. Furthermore, in every single case, the LOB did not change whatsoever between orders.

Yet another explanation, the "sounding-out hypothesis", may be a market-compliant strategy, but the European Securities and Markets Authority (ESMA) has already identified it as an area of particular concern.[10] According to this hypothesis, ROECC orders are a trading strategy in which a participant uses a buy (or sell) order to place the best bid (or ask) price in the LOB for a short period of time, the intention being to either have the order executed during this period with a latent but hidden market order on

8 The wider the current spread, the less likely it is that an order originating from the other side of the market will bridge the spread and be executed by the ROECC order. Therefore, the algorithm can set longer lifetimes for wider spreads.
9 Self-trading occurs when two almost identical algorithms in the market interact unintentionally, with each then responding to the other's activity in a kind of positive feedback loop and revolving, in a sense, around the other.
10 See the section entitled "Ping orders" in ESMA 2012/122, p 21.

the selling (or buying) side[11] or to use the offer to draw the other side into entering an executable new order. The order is cancelled if neither of the intended events occurs within a very short space of time. The extremely quick cancellation of the ROECC order combined with the minimum response time in the Eurex system usually make it practically impossible for other market participants to deliberately respond to such orders, leaving the first variant as the only possible explanation.[12] Judging by the characteristics of ROECC events outlined in this box, this particular theory is an interesting possible explanation.[13]

Another conceivable explanation is known as "quote stuffing" and is not a permitted practice. Quote stuffing is the deliberate

Relationship between ROECC order lifetime and bid-ask spread

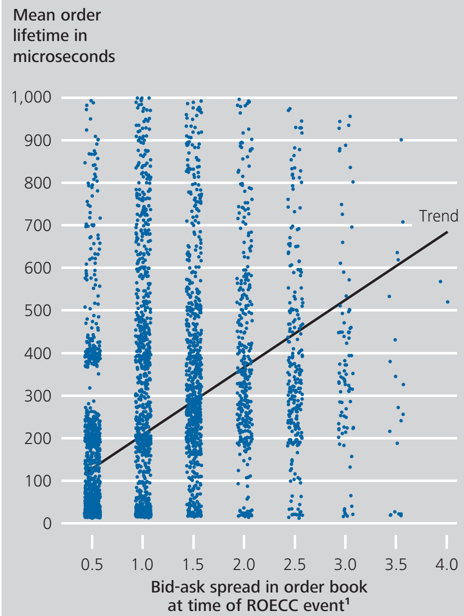

Sources: Eurex and Bundesbank calculations. **1** X-axis values dispersed slightly to aid visualisation. The genuine values are all 0.5; 1.0; 1.5 … 4.0.

Deutsche Bundesbank

entry and immediate cancellation of a large number of orders for a particular asset at a single trading venue.[14] One objective can be to temporarily slow down the exchange's trading system by generating a higher flow of orders. Perpetrators might then find it easier, say, to engage in profitable arbitrage because they could trade the asset affected by the marginal slowdown at a different trading venue at a more rewarding price. Given the low number of cycles in ROECC events (rarely more than a few dozen per second) and Eurex's server capacities (many thousand per second), this explanation can be ruled out.[15]

Another objective is to disrupt and mislead other algorithmic market participants. From the perspective of other trading algorithms, a vigorous strategy of sending and immediately cancelling a large number of orders generates significant data flows which they need to process, tying up their IT capacities. In this scenario, an ROECC issuer might be

11 If full execution is currently only possible at a very unfavourable price (in what is known as the "market order matching range"), many market models will briefly transform a market order into a kind of "waiting position", ready to be executed at a later point in time when conditions have improved for the party placing the order. Such waiting market orders are hidden from other market participants and represent a kind of latent liquidity – though they can interact immediately with new incoming orders.
12 The minimum response time in the Eurex trading system was around 180 microseconds between 2013 and 2015. Thus, other market participants will briefly see this order but have no opportunity to act on it.
13 The academic literature likewise contains evidence supporting this explanation. See J Hasbrouck and G Saar (2007), Technology and liquidity provision: the blurring of traditional definitions, Working Paper, Stern School of Business.
14 The law as it stood until June 2016 defined quote stuffing as a market manipulation strategy (within the meaning of section 20a (1) sentence 1 number 2 of the German Securities Trading Act (WpHG) because such behaviour could be conducive to disrupting or delaying the proper functioning of exchange operators' trading systems (section 3 (1) number 4a of the German Market Manipulation Definition Regulation (MaKonV)) or make it difficult for other market participants to identify the current buy and sell orders in the trading system (section 3 (1) number 4b of the MaKonV). See P Kasiske (2014), Marktmissbräuchliche Strategien im Hochfrequenzhandel, Zeitschrift für Wirtschafts- und Bankrecht, 68 (41), pp 1933-1939. From July 2016, section 20a of the Securities Trading Act and the Market Manipulation Definition Regulation were replaced by the EU Market Abuse Regulation without changing any aspects of the market conformity assessment.
15 This statement is based on information provided by Eurex representatives familiar with the IT architecture.

able to briefly slow down the algorithms of rival market participants and thus gain a decisive edge. The ROECC orders which supposedly suggest a misleading intention to buy or sell a given asset also have the effect of distorting the actual supply and demand situation, which again complicates matters for other market participants. This approach is likewise a possible explanation for ROECC events. One major drawback of this hypothesis, however, is that in a great many cases, an ROECC event is not actually immediately followed by a trade. Overall, it can be said that the observed cases cannot be attributed beyond a shadow of a doubt to quote stuffing behaviour, nor is it possible to identify any negative impact on market integrity.

A conclusive investigation into the intentions that lie behind ROECC events still needs to be conducted. What the analysis does illustrate, though, is that interesting phenomena occur in the space of extremely short timeframes during trading that are not perceptible to a human trader's eye, about which much still remains unknown. This shows, then, that research based on trading data in the sub-second range can make an important contribution to the understanding of today's markets.

… but is sensitive to foreseeable volatility events

The withdrawal of the passive HFT traders observed in the run-up to the release of the US labour market data can be interpreted in terms of their liquidity-enhancing strategies having a marked opportunistic character and reacting very sensitively to anticipated volatility events. For the days under consideration here, the hypothesis that passive HFT traders make a continuing significant contribution to liquidity provision even if there is strong expected volatility cannot be confirmed, at least for the period around the publication of the US labour market data.[42] At the time of the announcement at 14.30, the small amount of remaining liquidity is provided almost exclusively by NHFT traders. It should nevertheless also be noted that these are rare and only short-term withdrawals which appear to represent more of an exception than a rule.

Summary and outlook

While active HFT traders have a stronger presence in periods of market volatility …

Based on a new and granular database, the importance of HFT for market stability and integrity in trading in DAX and Bund futures is highlighted. These are the two most liquid German investment instruments in which HFT accounts for a significant share of trading activity.

The results suggest that HFT traders participate more strongly in trading as active market players in the Bund future especially in times of higher volatility. This applies both in periods of unexpectedly occurring volatility and in the run-up to expected price fluctuations. Especially following the announcement of important news, HFT traders can exploit their speed advantage; in doing so, they improve price discovery on very small time scales, although they

42 This stands in contrast to numerous traditional market makers, as binding rules on many trading platforms oblige them to provide liquidity in the order book even in adverse market conditions.

Share of HFT liquidity 500 seconds before and after publication of US labour market data*

Percentage share of liquidity

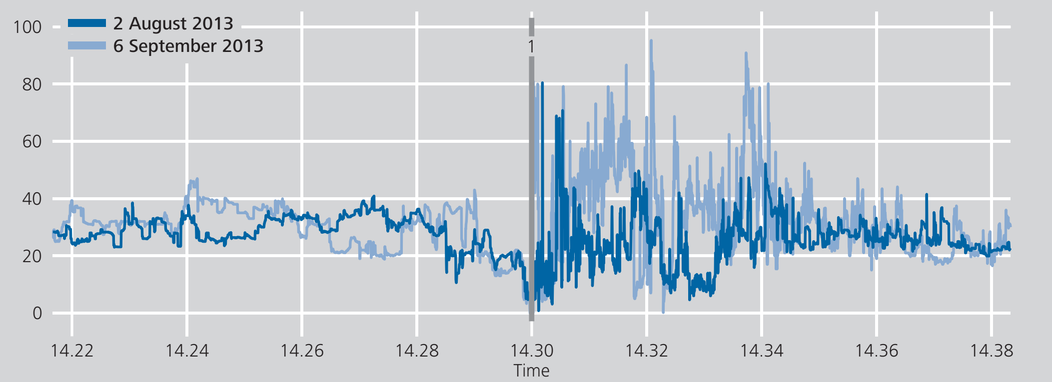

Sources: Eurex and Bundesbank calculations. * Up to 5 DAX points away from the best bid-ask price. **1** Publication of US labour market data.
Deutsche Bundesbank

also contribute to an above-average extent to short-term (excess) volatility.

… passive HFT traders reduce their supply of liquidity

At the same time, the results point to passive HFT participants often withdrawing in periods of market volatility and reducing their supply of liquidity. Taken together, the differing behaviours of active and passive HFT traders suggest a heightened risk of episodes of excessive short-term volatility which could provoke market turmoil, including flash events.

Analysis of the order book reveals the importance and temporal dynamics of HFT liquidity

Reconstructing the DAX future order book reveals the important role played by passive HFT traders in a good provision of liquidity that is stable over time. However, the results for the DAX future also confirm the findings for the Bund future that the liquidity supplied by HFT decreases significantly in times when important news is announced. A further finding is the existence of a rare, but noticeable pattern of behaviour in deletions of orders on very small time scales. Although it is not possible to pinpoint the cause of this anomaly, the phenomenon does illustrate the part that an in-depth study of highly granular trading date can play in helping regulators gain a better understanding of the market. Taking an overall view, it should be borne in mind that the obtained results relate solely to the futures markets for the DAX and Bunds under study here and can vary widely depending on the market segment and the share of HFT.

Rare and unusual patterns of activity on very small time scales

Greater data transparency is essential for further studies

Future studies by central banks, regulatory authorities and academic institutions might benefit from easier access to similarly granular data.[43] Enhanced transparency would also make it significantly easier to analyse new variants of intermediation in the capital markets in a timely and accurate manner. Understanding of how modern electronic trading works could be substantially improved in this way.

Results can assist debate on HFT regulation

The empirical results also underline possible points of approach in the regulatory debate about HFT. First, they demonstrate how important it is to implement incentive mechanisms so that passive HFT market makers maintain the provision of liquidity even in periods of heightened stress in the market.[44]

43 One example of future possibilities for analysis would be the use of the trader ID of individual transactions. A trader ID is an algorithm, contained in the dataset of every single order, which identifies the market participant placing the order. This makes it possible to answer the question of who has placed an order.
44 Article 17 (3) of the EU directive on markets in financial instruments (MiFID II) stipulates that an investment firm that engages in algorithmic trading (not just HFT traders) to pursue a market making strategy shall carry out this market making continuously during a specified proportion of the trading venue's trading hours, except under exceptional circumstances.

Second, the results suggest that active HFT traders, owing to their speed advantages in responding to the publication of important news, contribute to an excessive temporary volatility rather than to an informative one. This might permanently discourage slower market participants from providing adequate liquidity in such periods. Various instruments to counteract this problem are under discussion. These include a switch from continuous, steady trading to a discrete sequence of auctions (frequent batch auctions) and the introduction of a minimum time lag in the execution time of the orders of all market participants.[45,46] Both measures have in common that the resulting delay can restrict all the market players in their ability to respond by fractions of a second. Slower passive liquidity-providing participants would thus have a better chance of adjusting their orders to current market conditions.[47] This would partly offset the competitive disadvantages of slower market participants, about which there is much criticism in the public debate, without perceptibly impairing technological progress on the trading platforms. Not least, it would reduce the incentives for what is – in terms of the economic benefit – a dubious technological "arms race" on the trading platforms.

45 In a batch auction, a large number of incoming trading orders are combined and executed together at brief intervals in the form of an auction. This can reduce the speed advantage of HFT traders. See E Budish, P Cramton und J Shim (2015), The high-frequency trading arms race: frequent batch auctions as a market design response, Quarterly Journal of Economics, 130(4).
46 Since spring 2016, the electronic trading platform IEX Group has been delaying trading in shares on Wall Street by 350 microseconds.
47 See T Foucault (2016), op cit.