DE GRUYTER OLDENBOURG Review of Economics 2020; 71(3): 169–195

Antonio Sánchez Serrano*

High-Frequency Trading and Systemic Risk: A Structured Review of Findings and Policies

https://doi.org/10.1515/roe-2020-0028
Received September 7, 2020; accepted December 4, 2020

Abstract: A wider use of technology has contributed to the rapid growth of trading in stock markets in the last decades, resulting in an increase in the number of participants and a sharp decline in the price of information. High-frequency trading could be seen as a manifestation of this development. A review of the main findings in the academic literature leads to the identification of four main systemic vulnerabilities related to high-frequency trading: (i) adverse selection in orders, with the potential of crowding-out non-HFT market makers in times of stress; (ii) correlation of positions and herd behaviour; (iii) market power that, via technological costs, may impose barriers to entry; and (iv) negative contribution, in some circumstances, to price discovery. The first vulnerability could create systemic risk and several scholars have discussed the introduction of a limit in the speed of trading to address it. This could also contribute to reduce market power of high-frequency traders and over-investment in information technologies. Despite intense research efforts, further data and research is still needed to better understand these vulnerabilities and the adequacy of policies to address them.

Keywords: high frequency trading, systemic risk, market-making, markets microstructure

1 Introduction and Motivation

Stock exchanges were created as places to trade financial assets at a time when information was scarce and expensive, making it difficult to connect savers and entrepreneurs. Central exchanges, understood as places where securities were traded, appeared first in coffee houses of Amsterdam, London, New York and Paris in the 17th century, mostly for the provision of funds to overseas expeditions. For several centuries, only few investors had the appropriate resources to take part in them, generating a limited amount of transactions. The evolution of stock markets

***Corresponding author: Antonio Sánchez Serrano**, European Systemic Risk Board, Frankfurt, Germany, E-mail: antonio.sanchez@esrb.europa.eu. https://orcid.org/0000-0001-7670-0802

has intensified with the development of information technologies (IT) in the last 30 years, in a continuous and subtle process rather than a one-off event. As a result, it has been observed that (i) the number of participants in financial markets has exploded, as information technologies and the development of finance in general have enabled more economic agents to acquire the capabilities to trade,[1] and (ii) the price of information has plummeted, making it available to interested parties almost immediately at a very low cost.

Considering long-term trends in stock markets, an increase in the reported trades and a reduction in the average size of each trade has been observed (see, for example, Hendershott et al. 2011; US Securities and Exchange Commission 2010). As shown by Figure 1, in the 1980s and 1990s, transactions had an average size over 1500 dollars, while an average transaction involves less than 300 dollars since 2007.

Stock exchanges have also seen a substantial increase in liquidity, measured through turnover ratios (Levine and Zervos 1996), as the growth in the number of traded shares (numerator of the turnover ratio) has grown more than the market capitalisation (denominator of the turnover ratio). Besides, there is widespread consensus in the academic literature on the continuous decrease in bid-ask

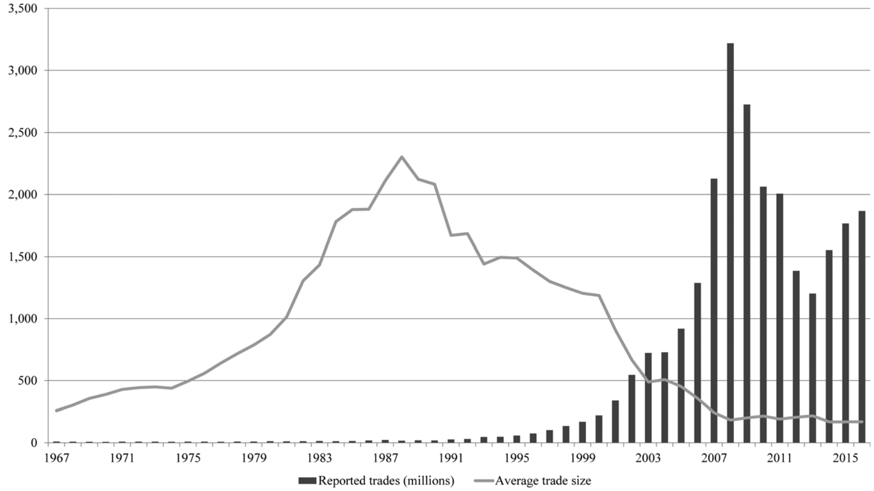

Figure 1: Reported trades and average trade size, NYSE markets, 1967–2016.
Source: NYSE and author's elaboration. Notes: Average trade size is calculated as the reported share volume divided by the number of reported trades.

1 For example, more than half of households in the US own shares.

spreads since the 1980s (Angel et al. 2011; Hendershott et al. 2011; Jones 2013). Technology, which started to get into stock markets in the 1980s, is supposed to have played a determinant role on that decrease of bid-ask spreads.

In this paper, we review the existing academic literature on high-frequency trading (HFT)[2] and, based on it, derive implications for systemic risk. The "flash crash" of 2010 as well other similar episodes of lesser magnitude have been associated with the growing role that HFT play in financial markets, with the potential to hamper their stability and smooth functioning. Based on the findings of the academic literature, we identify the main vulnerabilities through which HFT can contribute to systemic risk and we discuss how policies taken by regulatory and supervisory authorities can prevent and mitigate the materialisation of these vulnerabilities.

There is not a common worldwide definition of HFT, although the existing ones tend to agree in the identification of the key characteristics of HFT: (i) automated processes for trading, (ii) high speed in the submission of orders and in the process of incoming information, and (iii) submission of high numbers of orders and/or quotes (Chung and Lee 2016; European Parliament and Council 2014; O'Hara 2015; US Securities and Exchange Commission 2010). It is important to note that, conceptually, HFT is a part of automated trading, but not all automated trading necessarily fits into the definition of HFT. Besides, these definitions do not provide a precise quantitative definition of "high speed", although HFT are often seen to trade at intervals measured in milliseconds, a period of time beyond human comprehension.

It is difficult to accurately derive causality between HFT and any of the structural changes in financial markets referred above. However, the evolution of some key indicators in financial markets suggest the existence of such interaction. To start with, the increase in the turnover ratio and the high number of trades and contracts reflect the way HFT operates, which is based on the submission of a large number of orders during a trading day. Besides, the reduction in the average size of the transactions and the use of dark pools might be the reaction of some market participants to HFT, by either reducing the size of their intended transaction (to avoid falling under the radar of HFT strategies) or by recurring to more opaque venues with restricted access. Finally, lower bid-ask spreads can be considered to be influenced by the exploitation of arbitrage opportunities by HFT.

At this stage, it is important to consider how systemic risk in financial markets is defined. We leave aside more generic definitions of systemic risk (Billio et al. 2012; Bisias et al. 2012; De Bandt and Hartmann 2000; Smaga 2014) and we focus on

2 Throughout the text, the acronym HFT is used indistinctively to denote "High-Frequency Trading" as an activity as well as "High-Frequency Traders" as the agents undertaking HFT. The context where the acronym appears should allow its correct interpretation.

those considering systemic risk in financial markets. In this regard, Poledna et al. (2015) define systemic risk in financial markets as the risk that a significant fraction of the financial system can no longer perform its function as a credit provider and collapses. Pagano et al. (2019) define systemic risk in financial markets as the risk of a disruption in the financial system that leads to widespread instability in asset prices, extensive and material losses by investors and financial intermediaries, and possibly the collapse of important financial intermediaries. Key elements of both definitions are going to be considered in this paper, as they are complementary: Pagano et al. (2019) focus on how systemic risk may materialise (asset price instability, material losses to participants), and Poledna et al. (2015) refer to the consequences of systemic risk, leading to market collapse and disruption in the key function as credit providers (or connecting savings with investment opportunities). Both definitions imply that situations where investors face substantial losses do not automatically lead to systemic risk.

Furthermore, Poledna et al. (2015) identify two mechanisms through which systemic risk in financial markets can emerge: through the synchronisation of agents' behaviour (fire sales, margin calls, herding), or through their interconnectedness. In this paper, we will not consider interconnectedness as a source of systemic risk in financial markets, as that would probably merit a separate piece of research. In the context of HFT, Kumiega et al. (2016) state that systemic risk may emerge from the synchronized behaviour of large algorithmic traders, the introduction of unintended order messages and transactions by out of control algorithms, or a cybersecurity breach.

The impact of HFT on financial markets, in general, and its interaction with systemic risk, in particular, has been subject to an intense discussion in the academic literature. The main contribution of this paper to the existing literature refers to its staged approach: it starts with the main theoretical models, then presents the existing empirical evidence grouped into four main vulnerabilities and it finishes with a discussion on possible policy actions to address these vulnerabilities. The main theoretical models produced in the academic literature, focusing on specific aspects of the relationship between HFT and systemic risk, are presented first, in order to later on document the existing empirical evidence and construct a narrative on the interaction between HFT and systemic risk and discuss the effectiveness of different policies available to regulators. Several reviews on the academic literature on HFT cover the growing literature on the topic (Jones 2013; O'Hara 2015; Virgilio 2019). In its attempt to classify the vulnerabilities arisen from HFT on the basis of existing empirical evidence, this paper is closer to Menkveld (2016), Foucault and Moinas (2018), and Linton and Mahmoodzadeh (2018). In what regards its policy discussion, this paper can be placed together with Goldstein et al. (2014) and Chung and Lee (2016), as it bridges the main academic findings and the available policy measures to address them. The recent implementation of some of these policies is allowing now for the build-up of academic evidence on their effectiveness, some of which is also referred here.

Under that setting, our paper finds that, under certain circumstances, HFT may have implications for systemic risk, as the smooth functioning of financial markets could be affected. Among these implications, we consider, although not limited to, a potential amplification of the flight of liquidity in times of financial turmoil, certain herd behaviour and correlation of trading strategies, and the deviation of prices from their fundamental value on a continuous basis (as also outlined by Foucault and Moinas 2018, and Linton and Mahmoodzadeh 2018). At this stage, it is worth noting that these vulnerabilities have already been present in financial markets before HFT, so they should not univocally associated with HFT. For example, adverse selection is typically found to arise in financial markets with informed and uninformed traders, regardless of their speed of trading. However, HFT may exacerbate these vulnerabilities and the speed, depth and frequency with which they may materialise.

This paper is organised as follows. A survey of the relevant academic literature on HFT is provided in Section 2, and Section 3 outlines the main identified vulnerabilities derived from HFT. Section 4 discusses policies to address these vulnerabilities from a systemic risk point of view and Section 5 concludes.

2 Literature Review on HFT

The attention of the public at large towards HFT was caught by the “flash crash” of May 2010. In the academic domain, a rich stream of literature on HFT has developed in the last years, aimed at shedding light on the impact of HFT on financial markets. The following paragraphs provide a survey of the relevant academic literature on HFT organised around two main areas: (i) theoretical models on HFT, and (ii) the role of HFT in the “flash crash” of 2010. In addition to these two areas, there is a plethora of empirical academic work on HFT, which will be referred in the next section, where individual vulnerabilities are discussed in further detail. The authors and papers cited here do not pretend to be an exhaustive list of all the research in this field, but rather to present main arguments and lines of thought.

2.1 Theoretical Models on HFT

To start with, there have been an important number of attempts to model HFT, on a theoretical basis, and its impact on societal welfare.

Hoffmann (2014) develops a stylised model of trading in a limit order market where market participants can trade at two different speeds (slow and fast). In a comparison with an equilibrium with only one speed of trading, he finds that HFT have two opposing effects: (i) their ability to revise quotes after news arrivals reduces the existing inefficiency (increasing trading), but (ii) the speed advantage of HFT generates a new type of inefficiency, as other market participants tend to strategically submit limit orders with a lower execution probability (reducing trading). Additionally, there is shift in market power between HFT and other market participants, with the latter always worse-off in equilibrium. So in terms of welfare, the introduction of HFT reduces it when compared with a hypothetical setting with only one trading speed. Finally, he finds that HFT are more likely to act as market-makers in equilibrium.

Biais et al. (2015) develop a model of trading at different speeds, which considers also general welfare and allows for policy analysis. They point towards an overinvestment in technology by HFT, derived from the fact that they do not internalise the externalities they produce for slower investors (adverse selection). In terms of policies, their model suggests that welfare is maximised when slow and fast traders co-exist in a single market and there is a Pigouvian tax on investment in technology.

Foucault et al. (2016) propose a model of trading on news where they look at the optimal trading strategies of an informed fast speculator (who can trade ahead of incoming news) and of slower investors. They find that the fast speculator's trades represent a larger fraction of trading volume and are more correlated with short-run price changes. Fast speculators make a large proportion of their profits from trading on long-term price changes. Besides, both types of investors decrease their expected profit with news informativeness, but at a lower rate in the case of fast speculators. So, the relative gain of fast speculators increases with news informativeness. They also predict that the profitability of directional HFTs should be inversely related to news informativeness, but, at the same time, stocks with more informative news are more likely to attract directional HFTs.

In the model of Foucault et al. (2017), arbitrage opportunities arise when prices adjust to new information with a lag. These arbitrage opportunities are labelled as toxic because they expose dealers to the risk of trading at stale quotes. Their model also predicts that a technological change making arbitrageurs relatively faster should reduce the duration of arbitrage opportunities. Using data on triangular arbitrage, they find that illiquidity is higher on days when the fraction of toxic arbitrage opportunities and arbitrageurs' relative speed are higher (usually at a timescale of milliseconds). That would be consistent with the idea that price efficiency gains of HFT come at the cost of increased adverse selection risk.

Bongaerts and van Achter (2016) develop a model to analyse the impact on market stability of the liquidity provision by HFT. In their model, orders are limited and there is informed trading and endogenous entry and exit, with liquidity provided by both traditional market-makers and HFT. The latter trade at a higher speed through the use of higher information processing technologies. They find that speed technology leads to an efficient allocation of resources and maximises liquidity while information processing technologies can generate minor issues related to information asymmetries. However, the combination of the two can lead to inefficient speed technology, oligopolistic rents and increased issues related to information asymmetries.

Jain et al. (2016) use the launch of Arrowhead by the Tokyo Stock Exchange as a quasi-natural experiment to gauge the impact of high-frequency quoting (HFQ) and HFT on financial markets. While they find that the introduction of Arrowhead and the related increase in latency and speed of trading improved market liquidity, they are concerned about the increase in systemic risk, mainly in the form of autocorrelation and cross-correlation in the order flow (which can lead to an amplified propagation of shocks) and of more frequent "flash crashes" during periods of stressful market conditions. They use several measures of systemic risk developed in the academic literature and conclude that the introduction of Arrowhead led to increases in all of them.

From the perspective of trading venues, the model by Pagnotta and Philippon (2018) considers that trading venues invest in technology in order to attract investors who value speed and to whom they can charge higher fees. In principle, this competition among venues may be beneficial by increasing investor participation, trading volumes and efficiency. They also find that regulations that protect transaction prices (in a nutshell, investors have access to a single asset via two markets with a unique price) lead to greater fragmentation and faster speeds but might reduce efficiency. In terms of welfare, they conclude that a regulator would find it optimal to impose a minimum speed requirement but not a maximum speed limit. When considering price protection policies, their consequences in terms of welfare depend crucially on the impact on entry decisions of venues. When protection increases entry, the impact of price protection on welfare is typically positive. In some cases, the implicit subsidy embedded in price protection can allow entry by a slower venue, stimulate competition and result in higher investor participation and greater allocative efficiency. However, when price protection does not increase entry, it typically has a small negative impact on welfare derived from lower price competition and efficiency.

Keeping the focus on speed at exchange level, Menkveld and Zoican (2017) examine whether markets improve when trading platforms reduce their latency. They find that increasing speed can reduce liquidity. They identify two channels by which speed affects the bid-ask spread and the increase or decrease in liquidity depends on which channel dominates. The first channel implies that when an exchange increases its speed, then HFT can more quickly update their quotes, reducing the competitive spread. According to the second channel, when an exchange increases speed, HFT acting as market-makers become more exposed to speculative HFT every time new information is released. A faster exchange would then increase the spread. In terms of dominance, the first channel tends to dominate when the news rate is high, whereas the second channel is strong when there is intense competition among HFT.

More recently, Baldauf and Mollner (2020) consider the impact of HFT (and possible policy responses) in trading through the trade-off between liquidity and information production. Their model features random latency, multiple exchanges, and a single security that is traded by liquidity investors, information investors and HFT. They find that (i) faster trading leads HFT to be more successful at order anticipation (narrowing spreads and decreasing the cost of information); (ii) order anticipation pushes outcomes inside the frontier of this trade-off, representing an inefficiency of HFT; and (iii) this inefficiency stems from aggressive trading strategies. On the latter, they argue that some trading mechanisms can be designed to eliminate the inefficiency by preventing aggressive-side order anticipation. Among these mechanisms, they refer to non-cancellation delay mechanisms and frequent batch auctions.

Moving to the area of market design, Budish et al. (2015) argue that the development of HFT is a signal of flawed market design and they call for the setting of frequent auctions as a solution. Using millisecond-level data, they find that, at high-frequency time horizons, correlations completely break down, which introduces arbitrage opportunities, and that competition has not affected the size or frequency of the arbitrage opportunities. In the model they develop, they argue that these mechanical arbitrage opportunities stem from the continuous-time serial processing of orders. The rents which can be extracted from these arbitrage opportunities hinder the provision of liquidity and provide incentives to market participants to engage in a never-ending arms race for speed, which does not have any positive effect on welfare. As a policy response to this, they argue that the introduction of frequent auctions would transform competition in speed to a competition in prices, removing arbitrage opportunities and enhancing liquidity in markets.

2.2 Academic Literature on the "Flash Crash" of 6 May 2010

After a revision of some relevant theoretical work on HFT, it is useful to refer to the literature analysing the "flash crash" of 6 May 2010, as that event drew the attention of policymakers towards HFT. The “flash crash” lasted just 36 min and most of the losses were recovered soon afterwards. In its short lifespan, stock prices (as well as futures, options and ETFs) collapsed and rebounded very rapidly. For example, the Dow Jones Industrial Average dropped 998.5 points (about 9% from the opening), within minutes, and later recovered a large part of the loss. Some stock prices dropped to just 0.01 $ and the prices of others jumped to 100,000 $ in a matter of seconds.

Kirilenko et al. (2017) examine intraday trading between 3 May 2010 and 6 May 2010 in the E-mini S&P 500 stock index futures market, which was found to be at the origin of the “flash crash”. They study the second-by-second co-movement of inventory changes of (non-HFT) market makers and of HFT and, consistent with the findings in the literature, find that inventory changes of HFT positively co-move with contemporaneous (and lagged) prices but that co-movement is negative in the case of (non-HFT) market markers. Interestingly, such relationship between inventories and prices changed for (non-HFT) market makers during the “flash crash” in comparison with the previous days but did not change for HFT. In the case of HFT, it has been suggested that when certain traders are able to react faster than others to a signal, they enter into a kind of “latency arbitrage”, to the expense of slower traders. This behaviour of HFT would contrast with that of (non-HFT) market makers, mostly in line with the theoretical findings of Glosten and Milgrom (1985). In this vein, they find that, between 3 and 5 May, stock prices for which HFT were net buyers in a given second increased in the following second and remained higher for 20 s more, with a similar qualitative behaviour also observed on 6 May.

Easley et al. (2011) also analysed in detail the trading activities about various markets during 6 May 2010. They argue that the events on that day were the result of new dynamics in the financial markets and highlight the prominent role played by order toxicity in hampering liquidity provision. Order toxicity occurs when the order flow adversely selects non-HFT market makers who may be unaware that they are providing liquidity at a loss (Easley et al. 2012). If order toxicity rises unexpectedly, due to a larger portion of trades being originated from informed traders, then non-HFT market makers would face losses. The subsequent process of adjustment of non-HFT market makers, which now become consumers of liquidity, would add volatility to the markets, as was the case with the “flash crash”.

Lastly, Golub et al. (2012) analyse all the mini “flash crashes” in the US equity markets between 2006 and 2011. Together with the impact created by market regulation, they argue that the existing market structure with many trading venues is fragmenting overall liquidity, being “flash crashes” the consequence of it. HFT play a crucial role in this development as imbalances in liquidity across trading venues create opportunities for HFT to profit from them.

3 HFT and Systemic Risk: Potential Sources of Vulnerabilities

The debate on the potential systemic risk posed by HFT remains open and subject to certain controversy. The "flash crash" of 6 May 2010 acted as a "wake-up" call for regulators and brought HFT to the front-line of the regulatory agenda. However, the "flash crash" has not been echoed by episodes of comparable severity afterwards. At the same time, it has been documented that the profitability of HFTs seems to have abated, in a movement which has been interpreted as HFT reaching its physical limits (Kaya 2016; Meyer and Bullock 2017).

Even if HFT seem to have declined in the priority list of regulators and supervisors worldwide, it has changed the dynamics of financial markets, in a process which has been beneficial in many aspects, but which, on the other hand, may exacerbate existing vulnerabilities or create new ones in certain circumstances. This section identifies four vulnerabilities related to HFT, based on available empirical evidence.

3.1 Adverse Selection in Orders and Market-Making

The first vulnerability refers to adverse selection in orders, understood as a situation where market participation is affected by asymmetric information with HFT acting as "informed" traders (Aoyagi 2018; Biais and Woolley 2012; Biais et al. 2015; Brogaard et al. 2014; Carrion 2013; Easley et al. 2011; Foucault et al. 2016, 2017).[3] The informational disadvantage of some market participants explains then the risk premium market makers charge, as they need to compensate that "information" risk (Glosten and Milgrom 1985). Following Haldane (2011) and O'Hara (2015), adverse selection in a HFT context needs not to be interpreted as access to a richer set of information on the fundamentals of an underlying security, but rather in terms of faster reactions to new pieces of information. The distinction should be made between fast and slow traders, not between informed and uninformed traders. Hence, the importance for HFT of being as fast as possible and the important investments they undertake in co-location, IT equipment, software and infrastructure.

3 Adverse selection can also be seen as a natural process of market efficiency, where uninformed traders (i.e. those who are not able to process public information) are slowly crowded out by informed traders.

Adverse selection becomes an issue of particular relevance when a market operates with HFT market-makers and non-HFT market makers, as the latter may be trading with a counterparty better informed (faster) than they are. There is evidence in the academic literature that HFT can also follow and even dominate market making strategies (Boehmer et al. 2018; Brogaard et al. 2015; Hoffmann 2014; Malceniece et al. 2019). In financial markets with a strong presence of HFT, this can lead to what Easley et al. (2012) define as order toxicity (Foucault et al. 2017, use the term "toxic arbitrage" with a similar definition). In a period of uncertainty about prices, which can be driven by multiple automatic transactions by HFT (like in the "flash crash" of 2010), the information risk that non-HFT market makers face impedes them to fulfil their expected role and generates substantial losses to them. Non-HFT market makers cannot keep pace with the incoming flow of information about price developments. Consequently, these non-HFT market makers cease to operate as market makers, because they are not able to keep speed with HFT (i.e. they are slow) and they would then be forced to charge a higher premium for the information risk they are facing (Haldane 2011; Menkveld and Zoican 2017). In such circumstances, the information embedded in prices decreases substantially and slower non-HFT market makers are crowded out from the market (Bongaerts and van Achter 2016; Huh 2014). Market making functions are left fully to HFT and their automatic trading. The ultimate consequence of this could be a dysfunctional and self-induced process of price formation, like the one observed on 6 May 2010. This mechanism was already described, under a different configuration of financial markets, by Glosten and Milgrom (1985), who also considered the self-fulfilling mechanism leading to wide bid-ask spreads and, eventually, to the closure of the market.

In an stylised way, the price (P) of a stock z at time $t+1$ can be assumed to be a function of the past observed price at t and of the change in the available information about that stock (I_t^z), with a third term component covering all other residual factors affecting the stock price (u_t^z):

$$P_{t+1}^z = P_t^z + \alpha\,\Delta\left(I_t^z\right) + u_t^z \tag{1}$$

With short time intervals, the difference between prices is expected to be insignificant, unless major news arise or uncertainty raises materially. The longer the time interval, though, the larger would be the influence of available information and other factors in the price of the stock. For HFT, the units in which time is measured are probably milliseconds, whereas for non-HFT market makers (and other types of investors) time can be measured in seconds. Despite this difference in the time scale, in calm times, the third component of Equation (1) above is expected to remain low, so the difference in the time scale used by HFT and non-HFT market makers does not substantially affect the setting of prices in the next period. However, in times of stress in financial markets, the information risk which non-HFT market makers face may become too large, derived from the second and third components of Equation (1) and from the fact that, just immediately before $t + 1$, the last price they observed, at time t (measured in seconds), may be a stale price, not representative of the real price. Therefore, they are not able to make a sound estimate of the price in the next period (P^z_{t+1}) and would need to charge a large information premium to compensate for the risk they are taking. Ultimately, they may decide to quit the market, since they know that faster traders can make them assume substantial losses. The decision of non-HFT market makers to abandon the market may impair the price discovery process (assuming that all market makers would qualify, in principle, as informed traders), which could be entirely in the hands of HFT market makers.

The findings in Chakrabarty et al. (2019) seem to confirm the link between adverse selection and HFT. In an analysis of the impact of the SEC's naked access ban, a policy that slowed down some of the most aggressive HFTs, they conclude that trading costs declined driven by reduced adverse selection cost. Similarly, Brogaard et al. (2018) find that HFTs typically provide liquidity during episodes of extreme price movements when only one stock is affected. In situations when several stocks are under extreme price movements, position risk of HFTs is larger than during normal times, triggering internal controls in the area of liquidity management, resulting in an excess of liquidity demand. Using the staggered entry of Chi-X in 12 European equity markets, Malceniece et al. (2019) find that HFTs provide less liquidity during times of market stress. Similar findings are reached by Huh (2014), who concludes that HFT provide less liquidity when markets are volatile, because information asymmetry tends to be higher in those times. On the other hand, Brogaard et al. (2015) find that fastest traders in the NASDAQ OMX market (those subscribing to optional speed upgrade) behave primarily as market makers, improving the liquidity provision to slow and faster traders. They explain these findings as the speed upgrade is used to improve inventory management rather than for trading on short-lived information.

3.2 Correlation and Herd Behaviour

A second identified vulnerability of HFT refers to the diversity in financial markets, touching upon correlation and herd behaviour. As a starting point, Chaboud et al. (2014) find evidence of correlation in HFT strategies in the FX markets, and Smith (2010) has documented a high level of correlation in transactions involving 1500 shares or less. Similarly, Lafarguette (2016) finds common statistical patterns in HFT quotes. Boehmer et al. (2018) examine a group of 31 HFT in Canadians markets and identify an important heterogeneity in HFT strategies. However, they also find that 78% of the trades made by HFT can be explained by just three common strategies, suggesting an important correlation in the order flow of HFT. Malceniece et al. (2019) argue that HFT increase the co-movement of stocks returns, attributing more than half of that increase in co-movement to correlated trading strategies by HFT. Intuitively, the optimisation of algorithms in HFT would remove the human component (often, comprising irrational behaviour) in trading, leading, in principle, towards more homogenous strategies and reactions to new events (Linton and O'Hara 2011). In essence, the growing importance of HFT would imply that factors closely linked to human behaviour would not have a large weight on financial markets.

The impact on systemic risk from this vulnerability can be twofold. First, starting with the more unlikely, a shock simultaneously affecting several HFT may have negative consequences for them in terms of solvency, given the fact that HFT usually have very thin capital positions, ultimately leading to widespread failures of HFT (Biais and Woolley 2012). Such a scenario could have undesired effects for systemic risk in terms of interconnectedness and contagion, as it would be difficult to identify who is exposed to the failed HFT. Actually, counterparty risk acquires a new dimension in a scenario where HFT absorb an important part of the transactions. However, on the other hand, overnight inventories of HFT tend to be small (Jovanovic and Menkveld 2010), so their counterparty risk in case of a failure should be manageable.

Second, in normal trading times, the fact that HFT dominate the "active" trading (with passive investors just following them) coupled with certain convergence amidst all HFT strategies (Haldane 2011) could in principle lead to over-reaction to new information, in particular if that information is unexpected or of particular importance (Brogaard et al. 2014; Caivano 2015). Returning to Equation (1), this over-reaction to new pieces of information would be reflected in significant variations in the term $\Delta(I_t^z)$ over a short time period. In other words, it could lead to financial markets where movements in prices have a shorter and more acute nature. Indeed, Jiang et al. (2014) find that HFT increase volatility in the US Treasury market during, shortly before and after the announcement of macroeconomic news.

The over-reaction to new pieces of information could ultimately lead to herd behaviour and higher probabilities of tail events or flash crashes (Biais and Foucault 2014; Jain et al. 2016). Menkveld and Yueshen (2018) conclude that a breakdown of cross-market arbitrage activity increases the fragility of markets and could result in price crashes, using the "flash crash" to support their statement. In relation to flash crashes, Johnson et al. (2013) have documented 18,520 crashes and spikes between 2006 and 2011 with durations below 1.5 s. They directly relate these events to the increasing importance of computers in trading.

3.3 Market Power and Barriers to Entry

The third vulnerability stemming from the expansion of HFT in financial markets refers to market power and barriers to entry. In the empirical literature, it has been observed that a small number of HFT usually account for a very large proportion of the orders posted into a stock exchange (Clark 2012; Hagströmer and Nordén 2013; Kaya 2016; Kirilenko et al. 2017). Most of these orders are ultimately cancelled, but they allow HFT to determine the inside quotes on a particular security (Biais and Woolley 2012). Financial market participants are then divided in two groups depending on their ability to trade fast (Cartlidge and Cliff 2012; Haldane 2011; Virgilio 2017). The first, and less numerous, group enjoys several advantages from their privileged position, while the slower group does not. Concerns derived from a two-tiered financial markets have been usually mentioned when discussing policies to address HFT (Bhupathi 2010; Hoffmann 2014). In principle, financial markets regulation must ensure that all financial market participants have equal access to information and prices. However, HFT may be changing this perception, as it operates at time ranges which are not accessible to all market participants (Johnson and Zhao 2012).

General economic theory would argue that if HFT is so beneficial, then slower traders would make an effort to become HFT and, ultimately, the privileges enjoyed by HFT would disappear. However, in reality, there are frictions in financial markets and the costly technology used by HFT can create a barrier to entry for new participants. For example, and simply referring to one aspect of the many fixed costs in HFT, the Hibernia Express transatlantic cable between London and New York, which has a tested latency of less than 58.95 ms, is expected to have costed more than 300 million USD (Hasbrouck 2016). Similarly, Laughlin et al. (2014) have estimated a cost, including operations and infrastructure, exceeding 500 million USD for the 3 ms improvement in communication between Chicago and New York.

To better understand the barrier to entry created by the necessary technology to use HFT strategies, it is useful to consider the concept of capability, used in engineering. Following Kumiega et al. (2014) and Van Vliet (2017), the capability of a single HFT strategy (C) would be defined as follows:

$$C = \frac{\mu_n - c}{3\sigma_n} \tag{2}$$

where μ_n is the average expected return per trade, σ_n is the standard deviation of expected returns per trade, and c is the average fixed cost necessary to operate as a HFT. Accordingly, HFT would operate if the capability ratio of Equation (2) is larger than one, although a value of 1.33 is usually set as a threshold in engineering literature (Kumiega et al. 2014). In an environment of increasing technological costs for HFT, higher returns per trade or lower volatility would be required in order to be able to keep pace with HFT. For potential market participants, Equation (2) implies that they should be able to design highly profitable strategies to be able to compensate for the elevated fixed costs.

From a different perspective, HFT would have strong interests in maintaining the technological advantage over other financial market participants and this could lead to certain over-investment in technology, as signalled by Ye et al. (2013) and Biais et al. (2015). Hirshleifer (1971) shows the existence of significant incentives to trade on private information. When applied to HFT, HFT may overinvest in technology (c in Equation (2) would be excessively large) in the search for additional profits from trading with private information, far from the social optimum (Hoffmann 2014; Ye et al. 2013). The extent to which the investment in technology by HFT reflects mere competition or an unproductive arms race ultimately depends on the assessment of the contribution of HFT to price discovery (Jones 2013).

3.4 Contribution to Price Discovery

Indeed, the fourth vulnerability arising from HFT refers to the degree its various strategies contribute to price discovery, raising ethical concerns in some circumstances. Price discovery is understood to be the process of determining the price of an asset in a market through the interactions of buyers and sellers. The role played by HFT in this process has been found to be positive by, among others, Carrion (2013), Brogaard et al. (2014) and Benos et al. (2017), while Lafarguette (2016) finds a positive relation between HFT and price discovery, but only in normal times.

More nuanced conclusions are reached by Weller (2018), as he finds that algorithmic trading increases price efficiency with respect to acquired information but it reduces the available information to which prices respond. In terms of social welfare, he argues that the net effect from informational channels is almost surely a loss. Similarly, according to the model by Jarrow and Potter (2012), HFT may lead to mispricing of stocks (understood as deviations from their fundamental value), derived from collective and independent responses from HFT to common signals. Similar findings are presented by the model of Yang and Zhu (2020), where HFT detect slower investors' order flow in period one and then compete with these orders in period two in order to exploit this information for their own benefits. They also conclude that this behaviour by HFT poses a risk to the market.

These findings support the argument that the answer to the question whether HFT contributes to the price discovery process depends much on the strategies followed by HFT themselves (Benos and Sagade 2016). Following the identification by the Dutch Authority for the Financial Markets (2010) and Miller and Shorter (2016), HFT strategies can be classified into three main groups:

- Market making strategies. In this situation, HFT sets the price in a venue of a stock which is also quoted on another venue. HFT would then gain the spread between the prices in the two venues. If this mechanism is expanded to several venues, HFT engage in what is called cluster trading. Option trading is a traditional field for this kind of strategies.
- Statistical arbitrage. This strategy is based on finding arbitrage opportunities using statistics. For example, if stock prices temporarily stop behaving as could be expected based on statistical assumptions, this can be used as a signal for action by HFT, as it is possible to make sound forecasts about where the price will end up.
- Aggressive strategies. Their underlying rationale is to be faster than the rest of the participants in the financial market, so these strategies depend mostly on having the fastest systems or the fastest connection to the trading venues. Some examples of aggressive strategies are (i) searching out limit orders of small and slower investors by placing immediate or cancel orders (so the investor will always pay the maximum price for an order, earning the HFT the difference), (ii) analysing how other algorithms work so that the HFT can carry out arbitrage on them, (iii) placing small gradual orders to slowly move the market and benefit from the ensuing volatility through an option contract previously arranged, or (iv) anticipating orders from large investors so that the HFT can trade ahead of it and earn a profit.[4]

The positive contribution of market making strategies to price discovery seems to be clear (Huh 2014; Menkveld 2013), since HFT exploits mispricing and allows new information to be incorporated into prices. However, concerns about the role of HFT in the price discovery process start to emerge regarding statistical arbitrage

4 These strategies are usually nicknamed by market participants, which usually talk about "stuffing", when an unwieldy number of orders are placed generating congestion and impairing market access for slow traders; "smoking", implying the placement of alluring quotes, attracting slow market orders, but rapidly revised to less generous terms, before the slow orders reach the market; or "spoofing", when the fast trader wants to buy, it first places a bid, and then large ominous limit sell orders, to scare slow traders into hitting this bid (Biais and Woolley 2012).

and, particularly, aggressive strategies (Baldauf and Mollner 2020; Menkveld and Zoican 2017; Ye et al. 2013). Aggressive strategies, in particular, can lead to situations where the price moves away from the fundamental value of the stock (Egginton et al. 2016; Manahov 2016). The purpose of aggressive strategies to provoke a price increase or decrease and to generate a profit from establishing a long or short position on that price would not seem to improve the price discovery process. For example, Egginton et al. (2016) argue that episodic spikes in quoting activity by HFT (stemming from low latency strategies) lead to decreased liquidity, higher trading cost and increased short term volatility.

In this vein, the Dodd–Frank Act forbid in 2010 certain aggressive strategies and some others have been contested for a long-time. It is on this basis that O'Hara (2015) has claimed that financial markets are currently faster, but also less fair and more complex.

4 Policies to Address Identified Vulnerabilities

Once the main vulnerabilities posed by HFT have been presented and summarised in Table 1 below, it is the purpose of this section to consider potential policies to address them. Consequently, the policies discussed below should not be read as an exhaustive list. Indeed, they concentrate on the vulnerabilities identified in the previous section, leaving aside others that are probably more relevant for micro-prudential market regulators and supervisors. At the same time, it is worth noting that these vulnerabilities (adverse selection, herd behaviour, barriers to entry, aggressive trading strategies) were already present in financial markets before HFT, so they should not automatically be associated with HFT. However, HFT may exacerbate these vulnerabilities and the speed, depth and frequency with which they may materialise.

A situation where non-HFT market makers abandon the market in the middle of an episode of financial turmoil, might potentially have systemic consequences derived from (i) the significant drain of liquidity in that market at the time when it is precisely needed the most, and (ii) the crucial role which HFT (and their automatic algorithms) would play in the process of price discovery. The "flash crash" of 6 May 2010 provides an example of such situation, even if it was short-lived and did not have severe consequences over a longer horizon. To avoid a scenario like this and based on Equation (1), one alternative could be to limit the speed of trading in the market (called "resting rules" by Haldane 2011), so that non-HFT market makers would not be continuously in a disadvantage, in terms of adverse selection, vis-à-vis HFT market-makers. Such policy would define a minimum value for t in Equation (1). At the same time, introducing a limit to the speed of trading could also

Table 1: Vulnerabilities related to HFT and systemic risk as identified in the academic literature.

	Adverse selection	Correlation/Herd behaviour	Market power	Price discovery[a]
Main references	Haldane (2011), Easley et al. (2011), Biais and Woolley (2012), Carrion (2013), Brogaard et al. (2014), Biais et al. (2015), Brogaard et al. (2015), Foucault et al. (2016), Bongaerts and van Achter (2016), Foucault et al. (2017), Menkveld and Zoican (2017), Aoyagi (2018), Brogaard et al. (2018), and Chakrabarty et al. (2019)	Linton and O'Hara (2011), Haldane (2011), Johnson et al. (2013), Biais and Foucault (2014), Brogaard et al. (2014), Chaboud et al. (2014), Jiang et al. (2014), Caivano (2015), Lafarguette (2016), Jain et al. (2016), Menkveld and Yueshen (2018), Boehmer et al. (2018), and Malceniece et al. (2019)	Bhupathi (2010), Haldane (2011), Cartlidge and Cliff (2012), Johnson and Zhao (2012), Ye et al. (2013) Hoffmann (2014), Kumiega et al. (2014), Biais et al. (2015), Van Vliet (2017), and Virgilio (2017)	Jarrow and Potter (2012), Ye et al. (2013), Manahov (2016), Egginton et al. (2016), Menkveld and Zoican (2017), Weller (2018), Yang and Zhu (2020), and Baldauf and Mollner (2020)
Consequences for a smooth functioning of financial markets	Episodes of market-wide asset price instability and growing illiquidity	Increased volatility, high probabilities of tail events or flash crashes	Overinvestment in technology, tiering of financial markets participants	Rent absorption from deviations of prices from fundamentals
Link to systemic risk	Yes	Not conclusive	Likely	Likely

[a] The vulnerability on price discovery refers only to aggressive HFT strategies.

address the vulnerabilities related to excessive investment in technology (it would not compensate to invest heavily in technology), constraining the *c* variable in Equation (2) and removing some of the current barriers to entry into the market. Ye et al. (2013) conclude that a restriction in the speed of trading would benefit all investors, including HFT.

The establishment of a limit to the speed of trading would make unnecessary other policies which have been proposed in this area, such as some firm commitment by all market makers to provide liquidity at all times (Haldane 2011; Jones 2013). These policies would be confronted, at least, with the fundamental challenge of how to force some market participants to remain active in the market at a time where it would be better for them to exit the market.

Existing academic evidence on the introduction of a limit to the speed of trading is mixed. Shkilko and Sokolov (2016) empirically show that when weather-related events between Chicago and New York remove the speed advantages of HFT, trading costs and adverse selection are lower. They also show that a long-term removal of speed differentials would result in similar effects and would increase gains from trade. However, Pagnotta and Philippon (2018) explicitly argue that this policy measure would not optimise welfare, Huh (2014) considers that such policy could create a disadvantage for HFT market-makers and Brogaard et al. (2015) find a positive relation between trading speed and liquidity.

In the last years, certain stock exchanges have introduced measures that are similar, in spirit, to a limitation in the speed of trading. For example, since 2016, the SEC allows certain stock exchanges (including Investors Exchange, IEX, New York Stock Exchange's NYSE MKT and Nasdaq) to introduce intentional delays by implementing different versions of speed bumps. These speed bumps apply only to those traders who take liquidity (in other words, those trading against standing orders from market makers). In a first assessment of the effectiveness of these policies, Hu (2018) finds that price discovery improved, and trading costs decreased after IEX introduced intentional delays. Similarly, Chakrabarty et al. (2018) find that quote-to-trade ratios decreased after the implementation of the first speed bump in these markets and, more interestingly, that markets without planned speed bumps have lost market share, have reduced their investment in high-speed assets, the share prices of their stocks have generated lower returns, and their stocks have become more attractive for short sellers.

An interesting related policy has been proposed by Budish et al. (2015), Manahov (2016), and Baldauf and Mollner (2020), who suggest implementing an auction system, at very short time intervals, in financial markets, as a way to control the speed race by HFT. These auctions would occur at intervals below 1 s (Manahov 2016, proposes 30 ms), so they would not be perceived as a disruption by slower traders. The main difference with the limit to the speed of trading would be the entities affected by the policy: in this case the speed of the stock exchange is being slowed down. Imposing a limit to the speed of trading would seem, a priori, more difficult to monitor, as it would affect a very large number of institutions. Making the stock exchange to trade slower would imply, on the other hand, substantial costs of implementation for stock exchanges and could introduce certain complexity in the daily functioning of the exchange. The CBOE EBBO Periodic Auction Book has taken this idea forward and organises different auctions throughout the trading day.

An alternative to the implementation of a limit to the speed of trading could be to allow market authorities to implement these limits only in periods of financial stress, when it is more likely that liquidity may drain quickly. These measures are called circuit-breakers and have been in place in financial markets for some time. Circuit breakers are usually considered to be rather effective tools (see Kim and Yang 2004, for a rich survey of existing literature at that time, and European Securities and Markets Authority 2020, for a comprehensive and recent analysis) and to have contributed to avoiding another "flash crash", like the one on 6 May 2010 (Haldane 2011). Others, however, are more sceptical about their effectiveness (Kedrosky 2016; Subrahmanyam 1994). In a recent contribution, Hautsch and Horvath (2019) argue that the positive effects of circuit-breakers are offset by the negative impact on market quality and price stability, in the sense that circuit-breakers do not break volatility spirals and do not reduce price uncertainty. They identify a regulatory trade-off between the protective role of trading pauses and their adverse effects on market quality. From a conceptual point of view, while circuit breakers are, in principle, encoded into the engine of exchanges and are automatically activated, developments in financial markets might be quicker than the reaction of market authorities. In these occasions, HFT could potentially benefit from arbitrage opportunities arising, for example, from the circumstances calling for the activation of circuit-breakers. Bates (2015) has dramatically expressed it as "the police using bicycles to catch Ferraris".

To address adverse selection induced by HFT, it has also been proposed to establish a Tobin tax on orders posted to a stock exchange (Chung and Lee 2016). Biais and Foucault (2014) highlight that such taxes exist, in certain way, in France (where traders can cancel or modify up to 80% of their orders and then pay a tax of 1 basis points of the value of cancelled or modified orders) and in the Euronext markets, where a fee of 0.1 EUR is charged on each order when the ratio of orders-to-execution is above 100. These taxes may be seen as an effective and direct way of addressing the externalities created by HFT, but they do not seem to be the optimal solution.[5] In principle, a Tobin tax would also apply to the group of HFT pursuing strategies that are beneficial for financial markets. Actually, if the objective of introducing a Tobin tax is to dis-incentivise the negative contribution of certain HFT strategies to price discovery, it would be better to design policies specifically targeted at them, instead of reaching all HFT strategies with the same policy (Biais and Woolley 2012; Hoffmann 2014).

Actually, a second area of concern from a systemic risk perspective refers to those HFT strategies that hamper the price discovery process, as they could ultimately lead to material misalignments between asset prices and their fundamentals, with adverse consequences over the long-term. Here, the policy already in place in the US has been to ban these strategies and to effectively monitor that the ban is not circumvented. More analysis and research may be required to gain understanding on how the different HFT strategies affect price discovery, requiring that regulators follow very closely market developments in this area. Besides, the implementation of a ban on certain HFT strategies may affect the profitability of HFT, as, typically, aggressive strategies are the most profitable ones (Baron et al. 2019; Hendershott et al. 2011; von Muller 2012). At the aggregated level, though, these losses may be compensated by the benefits derived from a smooth functioning of financial markets, with a fluent price discovery process, satisfied investors and limited cases where prices deviate continuously from fundamentals.

There have also been voices calling for policies forcing HFT to have more skin-in-the-game, by, for example, setting a prudential framework or capital requirements for them (Biais and Woolley 2012). Such policies would mitigate possible widespread failures of HFT in case of unfavourable market dynamics (it must be recalled that HFT are typically highly leveraged and they tend to follow similar strategies). This is certainly a valid concern from microprudential purposes. Nonetheless, the impact of this policy on the most relevant vulnerabilities from a systemic risk angle (adverse selection, price discovery) seems limited.

The consideration of introducing a prudential framework for HFT can also respond to certain concerns on correlation of investments and herd behaviour. Confronted with an unexpected shock, the algorithms used by HFT could react similarly, with the potential to exacerbate the initial impact of the shock because a large portion of financial market participants would start from similar positions and would react similarly. In view of such scenario, it could be convenient to address this vulnerability with policies targeted particularly to HFT. However, before considering policies in such a sensitive area, it would be necessary to carry

5 In economic theory, the best way to address externalities is with incentives, not with direct regulation.

Table 2: Summary of policies to address identified systemic vulnerabilities related to HFT.

	Targeted vulnerability	References in academic literature	Other comments
Limit in speed of trading and speed bumps	Adverse selection Market power	Haldane (2011), Ye et al. (2013), Shkilko and Sokolov (2016), Hu (2018), and Chakrabarty et al. (2019) Not supportive: Pagnotta and Philippon (2018), Brogaard et al. (2015), and Huh (2014)	
Auctions at short-time intervals	Adverse selection Market power	Budish et al. (2015) and Manahov (2016): Baldauf and Mollner (2020)	
Circuit breakers	Adverse selection	Haldane (2011) and European Securities and Markets Authority (2020) Not fully supportive: Hautsch and Horvath (2019)	
Tax (or fees) on orders	Adverse selection Price discovery[a]	Biais and Foucault (2014) and Chung and Lee (2016) Not supportive: Hoffmann (2014)	More effective to ban certain HFT strategies (United States)
Prudential framework	Correlation/Herd behaviour	Biais and Woolley (2012)	Unclear link to systemic risk

[a] The vulnerability on price discovery refers only to aggressive HFT strategies.

out more research and analysis on the potential increased correlations and herd behaviour induced by HFT in financial markets.

5 Conclusions

Based on the main findings from the academic literature, we have defined four vulnerabilities created by HFT that can have systemic risk consequences and we have discussed some of the proposed policy responses to address these vulnerabilities (see Table 2).

While some of the policy responses seem more promising than others in addressing systemic risk posed by HFT, there is, certainly, a need to undertake more research to better understand how HFT and systemic risk interact and which regulatory actions can contribute to address the identified vulnerabilities. Particularly, as some of the initial policy responses to address risks posed by HFT have already been in place for several years, research efforts could concentrate on them, analysing their impact and whether they have met their regulatory objectives. Speed bumps would be a clear and promising example of this needed avenue of research.

Disclaimer: The views expressed in this paper are those of the author and do not necessarily represent the views of the European Systemic Risk Board (ESRB), any of its Member Institutions or the ESRB Secretariat. Comments by one anonymous referee are gratefully acknowledged. All remaining errors are mine.

References

Angel, J., Harris, L., and Spatt, C. (2011). Equity trading in the 21st century. Q. J. of Finance 1: 1–53.

Aoyagi, J. (2018) Coexistence of slow and fast markets and strategic speed choices by High-Frequency-Traders. Working Paper, Available at: https://ssrn.com/abstract=3221990.

Baldauf, M. and Mollner, J. (2020). High-frequency trading and market performance. J. Finance 75: 1495–1526.

Baron, M., Brogaard, J., and Kirilenko, A. (2019). Risk and return in high-frequency trading. J. Financ. Quant. Anal. 54: 993–1024.

Bates, J. (2015). Post flash crash, regulators still use bicycles to catch Ferraris. Traders Magazine Online News, 24 April 2015, Available at: http://www.tradersmagazine.com.

Benos, E. and Sagade, S. (2016). Price discovery and the cross-section of high-frequency trading. J. Financ. Mark. 30: 54–77.

Benos, E., Brugler, J., Hjalmarsson, E., and Zikes, F. (2017). Interactions among high-frequency traders. J. Financ. Quant. Anal. 52: 1375–1402.

Bhupathi, T. (2010). Technology's latest market manipulator – high frequency trading: the strategies, tools, risks, and responses. N. C. J. Law Technol. 11: 377–400.

Biais, B. and Foucault, T. (2014). HFT and market quality. Bankers, Markets and Investors, No. 128.

Biais, B. and Woolley, P. (2012). The flip side: high frequency trading. Financial World.

Biais, B., Foucault, T., and Moinas, S. (2015). Equilibrium fast trading. J. Financ. Econ. 116: 292–313.

Billio, M., Getmansky, M., Lo, A., and Pelizzon, L. (2012). Econometric measures of systemic risk in the finance and insurance sectors. J. Financ. Econ. 104: 535–559.

Bisias, D., Flood, M., Lo, A., and Valavanis, S. (2012). A survey of systemic risk analytics. Annu. Rev. Financ. 4: 255–296.

Boehmer, E., Li, D., and Saar, G. (2018). The competitive landscape of High-Frequency Trading firms. Rev. Financ. Stud. 31: 2227–2276.

Bongaerts, D. and van Achter, M. (2016). High-frequency trading and market stability. Working Paper, Available at: https://ssrn.com/abstract=2698702.

Brogaard, J., Hendershott, T.J., and Riordan, R. (2014). High-frequency trading and price discovery. Rev. Financ. Stud. 27: 2267–2306.

Brogaard, J., Hagströmer, B., Norden, L., and Riordan, R. (2015). Trading fast and slow: colocation and liquidity. Rev. Financ. Stud. 28: 3407–3443.

Brogaard, J., Carrion, A., Moyaert, T., Riordan, R., Shkilko, A., and Sokolov, K. (2018). High frequency trading and extreme price movements. J. Financ. Econ. 128: 253–265.

Budish, E., Cramton, P., and Shim, J. (2015). The high-frequency trading arms race: frequent batch auctions as a market design response. Q. J. Econ. 130: 1547–1621.

Caivano, V. (2015). The impact of high-frequency trading on volatility: evidence from the Italian market. Consob Quaderni di Finanza No. 80.

Carrion, A. (2013). Very fast money: high-frequency trading on Nasdaq. J. Financ. Mark. 16: 680–711.

Cartlidge, J. and Cliff, D. (2012). *Exploring the robot phase transition in experimental human-algorithmic markets*. Government Office for Science, DR25, London, Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/289051/12-1058-dr25-robot-phase-transition-in-experimental-human-algorithmic-markets.pdf.

Chaboud, A., Chiquoine, B., Hjalmarsson, E., and Vega, C. (2014). Rise of the machines: algorithmic trading in the foreign exchange market. J. Finance 69: 2045–2084.

Chakrabarty, B., Huang, J., and Jain, P.K. (2018). Exchange competition with levelled speed. Working Paper, Available at: https://ssrn.com/abstract=3197088.

Chakrabarty, B., Jain, P.K., Shkilko, A., and Sokolov, K. (2019). Unfiltered market access and liquidity: evidence from the SEC Rule 15c3-5. Manag. Sci. [forthcoming].

Chung, K. and Lee, A. (2016). High-frequency trading: review of the literature and regulatory initiatives around the world. Asia-Pacific J. Financ. Stud. 45: 7–33.

Clark, C. (2012). *How to keep markets safe in the era of high-speed trading. Essays on issues, No. 303*. Federal Reserve Bank of Chicago, Chicago, Available at: https://www.chicagofed.org/~/media/publications/chicago-fed-letter/2012/cfloctober2012-303-pdf.pdf.

De Bandt, O. and Hartmann, P. (2000). Systemic risk: a survey. ECB Working Paper Series No. 35.

Dutch Authority for the Financial Markets (2010). High frequency trading: the application of advanced trading technology in the European marketplace.

Easley, D., Lopez de Prado, M.M., and O'Hara, M. (2011). The microstructure of the 'flash crash': flow toxicity, liquidity crashes, and the probability of informed trading. J. Portfolio Manag. 37: 118–128.

Easley, D., Lopez de Prado, M.M., and O'Hara, M. (2012). Flow toxicity and liquidity in a high-frequency world. Rev. Financ. Stud. 25: 1457–1493.

Egginton, J.F., Van Ness, B., and Van Ness, R. (2016). Quote stuffing. Financ. Manag. 45: 583–608.

European Parliament and Council. (2014). Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU (recast).

European Securities and Markets Authority. (2020). Market impacts of circuit breakers – evidence from EU trading venues.

Foucault, T. and Moinas, S. (2018). Is trading fast dangerous? Toulouse School of Economics Working Paper Series, No. 18–881.

Foucault, T., Hombert, J., and Rosu, I (2016). News trading and speed. J. Finance 71: 335–382.

Foucault, T., Kozhan, R., and Tham, W.W. (2017). Toxic arbitrage. Rev. Financ. Stud. 30: 1053–1094.

Glosten, L. and Milgrom, P. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. J. Financ. Econ. 14: 71–100.

Golub, A., Keane, J., and Poon, S.H. (2012). High frequency trading and mini flash crashes. Working Paper, Available at: https://ssrn.com/abstract=2182097.

Goldstein, M.A., Kumar, P., and Graves, F.C. (2014). Computerized and high-frequency trading. Financ. Rev. 49: 177–202.

Hagströmer, B. and Nordén, L. (2013). The diversity of high frequency traders. J. Financ. Mark. 16: 741–770.

Haldane, A. (2011). The race to zero. In: *Speech at the international economic association sixteenth world congress, Beijing, China*. Bank of England, UK, Available at: https://www.bankofengland.co.uk/-/media/boe/files/speech/2011/the-race-to-zero-speech-by-andrew-haldane.pdf?la=en&hash=87FAD32EC83FCC5A04FEC1195C03F6761D415455.

Hasbrouck, J. (2016). Securities trading: principles and procedures. Manuscript, Available at: http://people.stern.nyu.edu/jhasbrou/TeachingMaterials/STPPms12a.pdf.

Hautsch, N. and Horvath, A. (2019). How effective are trading pauses? J. Financ. Econ. 131: 378–403.

Hendershott, T., Jones, C.M., and Menkveld, A. (2011). Does algorithmic trading improve liquidity? J. Finance 66: 1–33.

Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. Am. Econ. Rev. 61: 561–574.

Hoffmann, P. (2014). A dynamic limit order market with fast and slow traders. J. Financ. Econ. 113: 156–169.

Hu, E. (2018). Intentional access delays, market quality, and price discovery: evidence from IEX becoming an exchange. DERA Working Paper, Securities and Exchange Commission.

Huh, Y. (2014). Machines vs. machines: high frequency trading and hard information. Finance and economics discussion series 2014-33, Board of Governors of the Federal Reserve System.

Jain, P., Jain, P., and McInish, T. (2016). Does high-frequency trading increase systemic risk? J. Financ. Mark. 31: 1–24.

Jarrow, R. and Protter, P. (2012). A dysfunctional role of high frequency trading in electronic markets. Int. J. Theor. Appl. Finance 15, Available at: https://doi.org/10.1142/S0219024912500227.

Jiang, G., Lo, I., and Valente, G. (2014). High-frequency trading around macroeconomic news announcements: evidence from the US Treasury market. Bank of Canada Working Paper No. 56.

Johnson, N. and Zhao, G. (2012). *Brave New World: quantifying the new instabilities and risks arising in subsecond algorithmic trading*. Government Office for Science, DR27, London, Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/289045/12-1060-dr27-quantifying-instabilities-risks-in-subsecond-algorithmic-trading.pdf.

Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J., and Tivnan, B. (2013). Abrupt rise of new machine ecology beyond human response time. Sci. Rep. 3: 02627.

Jones, C.M. (2013). What do we know about high-frequency trading? Columbia Business School Research Paper No. 13–11.

Jovanovic, B. and Menkveld, A. (2010). Middlemen in limit-order markets. Working Paper, Available at: https://ssrn.com/abstract=1624329.

Kaya, O. (2016). High-frequency trading: reaching the limits. Global financial markets, Deutsche Bank Research.

Kedrosky, P. (2016). The dubious logic of stock-market circuit breakers. The New Yorker, 12 January 2016, Available at: https://www.newyorker.com.

Kim, Y. and Yang, J. (2004). What makes circuit breakers attractive to financial markets? A survey. Financ. Mark. Inst. Instrum. 13: 109–146.
Kirilenko, A., Kyle, A., Samadi, M., and Tuzun, T. (2017). The flash crash: high-frequency trading in an electronic market. J. Finance 72: 967–998.
Kumiega, A., Neururer, T., and van Vliet, B. (2014). Trading system capability. Quant. Finance 14: 383–392.
Kumiega, A., Sterijevski, G., and Van Vliet, B. (2016). Beyond the Flash Crash: systemic risk, reliability, and high frequency financial markets. J. Trading 11: 71–83.
Lafarguette, R. (2016). High-frequency trading, market volatility and regulation: the role of high-frequency quoting and dark pools. European Central Bank Macroprudential Bulletin, Issue 2, October.
Laughlin, G., Aguirre, A., and Grundfest, J. (2014). Information transmission between financial markets in Chicago and New York. Financ. Rev. 49: 283–312.
Levine, R. and Zervos, S. (1996). Stock market developments and long-run growth. World Bank Econ. Rev. 10: 323–339.
Linton, O. and Mahmoodzadeh, S. (2018). Implications of high-frequency trading for security markets. Annu. Rev. Econ. 10: 237–259.
Linton, O. and O'Hara, M. (2011). *The impact of computer trading on liquidity, price efficiency/ discovery and transaction costs*. Governement Office for Science, London, Available at: https://openaccess.city.ac.uk/id/eprint/13825/1/11-1276-the-future-of-computer-trading-in-financial-markets.pdf.
Malceniece, L., Malcenieks, K., and Putniņš, T. (2019). High frequency trading and comovement in financial markets. J. Financ. Econ. 134: 381–399.
Manahov, V. (2016). Front-running scalping strategies and market manipulation: why does high-frequency trading need stricter regulation? Financ. Rev. 51: 363–402.
Menkveld, A. (2013). High frequency trading and the new market makers. J. Financ. Mark. 16: 712–740.
Menkveld, A. (2016). The economics of high-frequency trading: taking stock. Annu. Rev. Financ. Econ. 8: 1–24.
Menkveld, A. and Yueshen, B. Z. (2018). The Flash Crash: a cautionary tale about highly fragmented markets. Manag. Sci. 65: 4470–4488.
Menkveld, A. and Zoican, M. (2017). Need for speed? Exchange latency and liquidity. Rev. Financ. Stud. 30: 1188–1228.
Meyer, G. and Bullock, N. (2017). Race for speed among algo traders hits peak. Financial Times, 30 March 2017, Available at: http://www.ft.com.
Miller, R. and Shorter, G. (2016). *High frequency trading: overview of recent developments*. Congressional Research Service.
O'Hara, M. (2015). High frequency market microstructure. J. Financ. Econ. 116: 257–270.
Pagano, M., Sánchez Serrano, A., and Zechner, J. (2019). Can ETFs contribute to systemic risk? Report of the ESRB Advisory Scientific Committee 9.
Pagnotta, E.S. and Philippon, T. (2018). Competing on speed. Econometrica 86: 1067–1115.
Poledna, S., Molina-Borboa, S., Martínez-Jaramillo, S., van der Leij, M., and Thurner, S. (2015). The multi-layer network nature of systemic risk and its implications for the costs of financial crises. J. Financ. Stabil. 20: 70–81.
Shkilko, A. and Sokolov, K. (2016). Every cloud has a silver lining: fast trading, microwave connectivity and trading costs. Working Paper, Available at: https://ssrn.com/abstract=2848562.

Smaga, P. (2014). The concept of systemic risk. Systemic Risk Council Special Paper No. 5.
Smith, R. (2010). Is high-frequency trading inducing changes in market microstructure and dynamics? Working Paper, Available at: https://ssrn.com/abstract=1632077.
Subrahmanyam, A. (1994). Circuit breakers and market volatility: a theoretical perspective. J. Finance 49: 237–254.
US Securities and Exchange Commission (2010). Concept release on equity market structure. Release No. 34-61358, Available at: https://www.sec.gov/rules/concept/2010/34-61358fr.pdf.
Van Vliet, B. (2017). Capability satisficing in high frequency trading. Res. Int. Bus. Finance 42: 509–521.
Virgilio, G. (2017). Is high-frequency trading tiering the financial markets? Res. Int. Bus. Finance 41: 158–171.
Virgilio, G. (2019). High-frequency trading: a literature review. Financ. Mark. Portfolio Manag. 33: 183–208.
Von Muller, C. (2012). Regulating high frequency trading: a micro-level analysis of spatial behavior, optimal choices, and Pareto-efficiency in high speed markets. University of St. Gallen Law & Economics Working Paper No. 2012-04.
Weller, B. (2018). Does algorithmic trading reduce information acquisition? Rev. Financ. Stud. 31: 2184–2226.
Yang, L. and Zhu, H. (2020). Back-running: seeking and hiding fundamental information in order flows. Rev. Financ. Stud. 33: 1484–1533.
Ye, M., Yao, C., and Gai, J. (2013). The externalities of high frequency trading. WBS Finance Group Research Paper.