



Market stability vs. market resilience: Regulatory policies experiments in an agent-based model with low- and high-frequency trading

Sandrine Jacob Leal, Mauro Napoletano

► To cite this version:

Sandrine Jacob Leal, Mauro Napoletano. Market stability vs. market resilience: Regulatory policies experiments in an agent-based model with low- and high-frequency trading. *Journal of Economic Behavior and Organization*, 2019, 157, pp.15-41. 10.1016/j.jebo.2017.04.013 . hal-03403589

HAL Id: hal-03403589

<https://sciencespo.hal.science/hal-03403589>

Submitted on 10 Dec 2021

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

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



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Market Stability vs. Market Resilience: Regulatory Policies Experiments in an Agent-Based Model with Low- and High-Frequency Trading

Abstract

We investigate the effects of a set of regulatory policies directed towards high-frequency trading (HFT) through an agent-based model of a limit order book able to generate flash crashes as the result of the interactions between low- and high-frequency traders. In particular, we study the impact of the imposition of minimum resting times, of circuit breakers, of cancellation fees and of transaction taxes on asset price volatility and on the occurrence and the duration of flash crashes. Monte-Carlo simulations reveal that HFT-targeted policies imply a trade-off between market stability and resilience. Indeed, we find that policies able to tackle volatility and flash crashes also hinder the market from quickly recovering after a crash. This result is mainly due to the dual role of HFT, as both a cause of flash crashes and a key player in the post-crash recovery.

Keywords: High-frequency trading, Flash crashes, Regulatory policies, Agent-based models, Limit order book, Market volatility.

JEL codes: G12, G01, C63.

1 Introduction

This paper studies the effects of a set of regulatory policies aimed at curbing the possible negative effects of high-frequency trading (HFT henceforth), and at reducing market volatility and the occurrence of flash crashes.

Over the past decade, HFT has sharply increased in US and European markets (*e.g.*, AMF, 2010; SEC, 2010; Lin, 2012, and references therein). HFT also represents a major challenge for regulatory authorities, partly because it encompasses a wide array of algorithmic trading strategies and partly because of the big uncertainty yet surrounding the net benefits it has for financial markets. Indeed, on the one hand, some studies have highlighted the benefits of HFT as a source of an almost continuous flow of liquidity (see *e.g.*, Brogaard, 2010; Menkveld, 2013). On the other hand, other works have pointed

to HFT as a source of higher volatility in markets and as a key driver in the generation of extreme events like flash crashes (see *e.g.*, SEC, 2010; Angel et al, 2011; Lin, 2012; Kirilenko and Lo, 2013), whose incidence has grown in the last decades (Golub et al, 2012; Johnson et al, 2012). The regulatory framework is complicated by the fact that - although many explanations have so far been proposed for flash crashes - no consensus has yet emerged about the fundamental causes of these extreme phenomena (see Haldane, 2014). Overall, the above-mentioned open issues call for a careful design of regulatory policies that could be effective in mitigating the negative effects of HFT and in hindering flash crashes (and/or dampening their impact on markets, see also CFTC and SEC, 2010; Prewitt, 2012; Haldane, 2014).

Earlier empirical and theoretical works have already attempted to study the effect of different sets of regulatory measures (*e.g.*, Westerhoff, 2008; Brewer et al, 2013; Vuorenmaa and Wang, 2014) and of some specific regulation policies such as trading halts (Westerhoff, 2003b, 2006; Yeh and Yang, 2010, 2013), financial transaction tax (Colliard and Hoffmann, 2013; Lavicka et al, 2014; Biais et al, 2015; Fricke and Lux, 2015), minimum resting times (Hayes et al, 2012), market design (Budish et al, 2013), cancellation fee (Friederich and Payne, 2015), position limits (Lee et al, 2011). However, these works have either not considered the role of HFT (*e.g.*, Westerhoff, 2003b, 2008), or they have treated flash crashes as resulting from an exogenous shock (*e.g.*, Brewer et al, 2013) or, finally, they have only focused on a very narrow set of policies (*e.g.*, Hayes et al, 2012; Vuorenmaa and Wang, 2014).

On these grounds, we contribute to the current debate about the regulatory responses to flash crashes and to the potential negative externalities of HFT by studying the impact of a set of policy measures in an agent-based model where flash crashes endogenously emerge out of the interplay between low- and high-frequency traders. The goal of this work is to shed some light on which policy measures are effective to curb volatility and the incidence of flash crashes and/or to fasten the process of price-recovery after a crash. To this end, we extend the ABM developed in Jacob Leal et al (2016) to allow for endogenous orders' cancellation by high-frequency (HF) traders, and we then use the model as a test-bed for a number of policy interventions directed towards HFT. This model is particularly well-suited and relevant in this case because, differently from existing works (*e.g.*, Brewer et al, 2013), it is able to endogenously generate flash crashes as the result of the interactions between low- and high-frequency traders. Moreover, compared to the existing literature we consider a broader set of policies, also of various nature. The list includes market design policies (circuit breakers) as well as command-and-control (minimum-resting times) and market-based (cancellation fees, financial transaction tax)

measures.

The model in Jacob Leal et al (2016) portrays a market wherein low-frequency (LF) agents trade a stock, switching between fundamentalist and chartist strategies according to strategies' profitability. HF agents differ from LF ones in many respects. First, unlike LF traders, activation of HF traders is not based on chronological time but it is *event-based*, *i.e.* it depends on specific market conditions (see *e.g.*, Easley et al, 2012). Second, HF agents adopt low-latency *directional* strategies that exploit the price and volume information released in the book by LF traders (cf. Aloud et al, 2012). Third, HF traders keep their positions open for very short periods of time and they pursue tight inventory management (Kirilenko et al, 2011). Lastly, HF deliberately cancel their orders based on expected profits (see Kirilenko et al, 2011; SEC, 2014, for a review of cancellation practices of HF traders).

After checking the ability of the model to reproduce the main stylized facts of financial markets, we run extensive Monte-Carlo experiments to test the effectiveness of policies which have been proposed and implemented both in Europe and in the US to curb HFT and to prevent flash crashes, namely the implementation of *i.*) trading halt facilities (both ex-post and ex-ante designs); *ii.*) minimum resting times; *iii.*) order cancellation fees; *iv.*) transaction tax. Computer simulations show that slowing down high-frequency traders, by preventing them from frequently and rapidly cancelling their orders, ought to the introduction of either minimum resting times or cancellation fees, has beneficial effects on market volatility and on the occurrence of flash crashes. Also discouraging HFT via the introduction of a financial transaction tax produces similar outcomes (although the magnitude of the effects is smaller). All these policies impose a speed limit on trading. Thus finding that they are valid tools to cope with volatility and the occurrence of flash crashes confirms the conjectures in Haldane (2014) about the need of tackling the “race to zero” of HF traders in order to improve financial stability. At the same time, we find that all these policies imply a longer duration of flash crashes, and thus a slower price recovery to normal levels. Furthermore, the results regarding the implementation of circuit breakers are mixed. We find that the introduction of an ex-ante circuit breaker markedly reduces price volatility and completely removes flash crashes. This is merely explained by the fact that this type of regulatory design precludes the huge price drop, source of the flash crash. In contrast, ex-post circuit breakers do not have any particular effect on market volatility, nor on the number of flash crashes. Moreover, they increase the duration of flash crashes.

Overall, our results indicate the presence of a fundamental trade-off characterizing HFT-targeted policies, namely one between market stability and market resilience. Poli-

cies that improve market stability - in terms of lower volatility and incidence of flash crashes - also imply a deterioration of market resilience - in terms of lower ability of the market price to quickly recover after a crash. This trade-off is explained by the dual role that HFT plays in the flash crash dynamics of our model. On the one hand, HFT is the source of flash crashes by occasionally creating large bid-ask spreads and concentrating orders on the sell side of the book. On the other, HFT plays a key role in the recovery from the crash by quickly restoring liquidity.

The paper is organized as follows. Section 2 describes the model. In Section 3, we present and discuss the simulation results in three steps. First, we assess the ability of the model to jointly reproduce some of the common stylized facts of stock returns, also detected at very high frequency. Second, we discuss the main features of flash crash dynamics in our model. Finally, we present the results concerning policy experiments. Section 4 contains the concluding remarks. Lastly, the Appendix at the end of the paper presents the results of some robustness analyses concerning the activation mode of high frequency traders and their pricing strategy.

2 The Model

We use the model, developed in Jacob Leal et al (2016), of a stock market populated by heterogeneous, boundedly-rational traders. Agents trade an asset for T periods and transactions are executed through a limit-order book (LOB) where the information about the type, the size and the price of all agents' orders is stored (see, for instance, Maslov, 2000; Zovko and Farmer, 2002; Avellaneda and Stoikov, 2008; Bartolozzi, 2010). The market is populated by two groups of agents depending on their trading frequency (*i.e.*, the average amount of time elapsed between two order placements), namely N_L low-frequency and N_H high-frequency traders ($N = N_L + N_H$). Although the number of agents in the two groups is kept fixed over the simulations, the proportions of low- and high-frequency traders change over time, as some agents may not be active in each trading session. Furthermore, agents of both types are different not only in terms of trading frequencies, but also in terms of strategies and activation rules. A detailed description of the behavior of LF and HF traders is provided in Sections 2.1 and 2.2.

In the model, a trading session is assumed to last one minute. At the beginning of each trading session t , active LF and HF agents know past market prices as well as past and current fundamental values of the traded asset. Based on the foregoing information set, each trading session t proceeds in the following way. First, each active LF trader submits a buy or sell order to the LOB market, specifying its size and its limit price.

Next, active HF traders start to place their limit orders in the book in a sequential manner and the size and the price of their orders are also displayed in the LOB. We assume that - once all orders have been inserted in the book - HF traders are able to compute the transactions that would take place given the existing book, their prices and, therefore, their expected profits.¹ They then use the computed expected profits to decide whether to confirm or to cancel their orders from the book.² We capture the last feature by assuming that the matching procedure takes place in two steps. First, a temporary matching session takes place. On the basis of this procedure, HF traders are able to compute expected profits and decide whether to confirm or cancel their orders. More precisely, in the temporary matching procedure HF traders are able to simultaneously compute all the orders that would be matched given the existing book, check the quantities and prices of all the orders that will be executed (including theirs) and compute the last transaction price. They then use this information to compute their expected profits. HF agents use the level of expected profits to decide whether to confirm/cancel their orders (see Section 2.2). Cancelled orders are removed from the book. After the temporary procedure is completed and HF orders have been confirmed or cancelled, the actual matching procedure occurs. In the actual matching, the actual trading session price (\bar{P}_t) is determined as the price of the last executed transaction in the trading session.³ LF and HF unexecuted orders rest in the book for the next trading sessions (γ^L and γ^H periods, respectively). Lastly, given \bar{P}_t , all agents compute their actual profits and LF agents update their strategy for the next trading session (see Section 2.1 below). Notice that the possibility that some orders are removed from the book depending on expected profits and before the actual matching process takes place implies that the participation of HF traders to market transactions is motivated by profit considerations. In addition, the possibility of orders' cancellation implies in general a difference between temporary and actual trading prices, as well as a difference between expected and actual traders' profits.

¹In particular, we assume that HF agents simultaneously compute transactions, prices and expected profits based on the same order book information.

²The assumption that HF orders are inserted after LF traders' ones and that HF agents are able to calculate expected profits before actual matching takes places are convenient ways of capturing one of the distinctive feature of HFT *i.e.*, their ability to rapidly process a large amount of information and to exploit low-latency strategies (see *e.g.*, Hasbrouck and Saar, 2013).

³The price of an executed contract is the average between the matched bid and ask quotes.

2.1 Low-Frequency Traders

In the market, there are $i = 1, \dots, N_L$ low-frequency agents who take short or long positions on the traded asset.⁴ The trading frequency of LF agents is based on *chronological* time, which is heterogeneous across LF agents and constant over time. In particular, each LF agents' trading speed is drawn from a truncated exponential distribution with mean θ and is bounded between θ_{min} and θ_{max} minutes.⁵

In line with most heterogeneous agents models of financial markets (*e.g.*, De Long et al, 1990; Lux and Marchesi, 2000; Farmer, 2002; Kirman and Teyssiere, 2002; Chiarella and He, 2003; Hommes et al, 2005; Westerhoff, 2008), LF agents determine the quantities bought or sold (*i.e.*, their orders) according to either a fundamentalist or a chartist (trend-following) strategy. In particular, we use the model specification proposed in Franke and Westerhoff (2009, 2011) which accounts for within-group heterogeneity through a noise term added to each of the demand functions and which received strong empirical support in Franke and Westerhoff (2012). More precisely, given the last two market prices \bar{P}_{t-1} and \bar{P}_{t-2} , orders under the chartist strategy ($D_{i,t}^c$) are determined as follows:

$$D_{i,t}^c = \alpha^c(\bar{P}_{t-1} - \bar{P}_{t-2}) + \epsilon_t^c, \quad (1)$$

where $0 < \alpha^c < 1$ and ϵ_t^c is an i.i.d. Gaussian stochastic variable with zero mean and σ^c standard deviation. If a LF agent follows a fundamentalist strategy, her orders ($D_{i,t}^f$) are equal to:

$$D_{i,t}^f = \alpha^f(F_t - \bar{P}_{t-1}) + \epsilon_t^f, \quad (2)$$

where $0 < \alpha^f < 1$ and ϵ_t^f is an i.i.d. Gaussian random variable with zero mean and σ^f standard deviation. The fundamental value of the asset F_t evolves according to a geometric random walk:

$$F_t = F_{t-1}(1 + \delta)(1 + y_t), \quad (3)$$

with i.i.d. $y_t \sim N(0, \sigma^y)$ and a constant term $\delta > 0$. After γ^L periods, unexecuted orders expire, *i.e.* they are automatically withdrawn from the book. Finally, the limit-order price of each LF trader is determined by:

$$P_{i,t} = \bar{P}_{t-1}(1 + \delta)(1 + z_{i,t}), \quad (4)$$

⁴We assume that LF traders are not able to employ low-latency trading since they process information and respond to market events with a scale that is equal or higher than the one of the trading session.

⁵See also Alfarano et al (2010) for a model with different time horizons in a setting different from ours.

where $\delta > 0$ and $z_{i,t}$ measures the number of ticks away from the last market price \bar{P}_{t-1} and it is drawn from a Gaussian distribution with zero mean and σ_z^{LF} standard deviation.

In each period, low-frequency traders can switch their strategies according to strategy's profitability. At the end of each trading session t , once the market price \bar{P}_t is determined, LF agent i computes her profits ($\pi_{i,t}^{st}$) under chartist ($st = c$) and fundamentalist ($st = f$) trading strategies as follows:

$$\pi_{i,t}^{st} = (\bar{P}_t - P_{i,t})D_{i,t}^{st}. \quad (5)$$

Following Brock and Hommes (1998), Westerhoff (2008), and Pellizzari and Westerhoff (2009), the probability that a LF trader will follow a chartist rule in the next period ($\Phi_{i,t}^c$) is given by:

$$\Phi_{i,t}^c = \frac{e^{\pi_{i,t}^c/\zeta}}{e^{\pi_{i,t}^c/\zeta} + e^{\pi_{i,t}^f/\zeta}}, \quad (6)$$

with a positive intensity of switching parameter ζ . Accordingly, the probability that LF agent i will use a fundamentalist strategy is equal to $\Phi_{i,t}^f = 1 - \Phi_{i,t}^c$.

2.2 High-Frequency Traders

As mentioned above, the market is also populated by $j = 1, \dots, N_H$ high-frequency agents who buy and sell the asset.⁶ Contrary to LF agents, HF traders employ low-latency technologies which enable them to place their orders with high speed. Moreover, HF agents differ from LF ones not only in terms of trading speed, but also in terms of activation and trading rules. In particular, contrary to LF strategies, which are based on chronological time, HF agents adopt trading rules framed in *event* time (see *e.g.*, Easley et al, 2012),⁷ *i.e.*, the activation of HF agents depends on the extent of the last price change observed in the market. As a consequence, HF agents' trading speed is *endogenous*. More specifically, each HF trader has a fixed price threshold Δx_j , drawn from a uniform distribution with support bounded between η_{min} and η_{max} . This threshold determines whether the agent will activate or not in the trading session t (see

⁶We assume that $N_H < N_L$. The proportion of HF agents vis-à-vis LF ones is in line with empirical evidence (Kirilenko et al, 2011; Paddrik et al, 2011).

⁷On the case for moving away from chronological time in modeling financial series see Mandelbrot and Taylor (1967); Clark (1973); Ané and Geman (2000).

Aloud et al, 2013, for a similar attempt in this direction):⁸

$$\left| \frac{\bar{P}_{t-1} - \bar{P}_{t-2}}{\bar{P}_{t-2}} \right| > \Delta x_j. \quad (7)$$

Active HF agents submit buy or sell limit orders with equal probability $p = 0.5$ (Maslov, 2000; Farmer et al, 2005).

HF traders adopt *directional* strategies that try to profit from the anticipation of price movements (see SEC, 2010; Aloud et al, 2012) and exploit the price and order information released by LF traders and by other HF traders (if any).

First, HF traders account for current order flows to determine their order size $D_{j,t}$ (Ranaldo, 2004). More specifically, HF traders' order size is drawn from a truncated Poisson distribution whose mean depends on volumes available in the sell-side (*buy-side*) of the LOB, if the order is a buy (*sell*) order.⁹ The ability of HF traders to adjust the volumes of their orders to the ones available in the book reflects their propensity to absorb LF agents' orders. Moreover, in order to account for empirical evidence indicating that HF traders do not accumulate large net positions (CFTC and SEC, 2010; Kirilenko et al, 2011), we add two additional constraints to HF order size. On the one hand, HF traders' net position is bounded between $\pm 3,000$. On the other hand, HF traders' buy (*sell*) orders are smaller than one quarter of the total volume present in the sell (*buy*) side of the book (see, for instance, Bartolozzi, 2010; Kirilenko et al, 2011; Paddrik et al, 2011).¹⁰

Second, HF traders account for current best ask and bid prices to set their order limit price. In particular, in the trading session t , each HF agent j trades near the best ask ($P_{j,t}^{ask}$) or bid ($P_{j,t}^{bid}$) available in the LOB at the moment when the agent places the order (see *e.g.*, Paddrik et al, 2011).¹¹ Accordingly, HF buyers and sellers' limit prices

⁸Note that the distribution support of the activation threshold is chosen so that it generates heterogeneous thresholds across HF traders. Moreover, the support of the uniform distribution used in the simulation exercises is quite large and includes also very small threshold values (see η_{min} and η_{max} in Table 1). It follows that in our simulations some HF traders may activate in presence of large as well as small price changes. Activation for large price changes is also in line with empirical accounts of recent flash crash events (e.g. Kirilenko et al, 2011), that indicate that HF traders were well active during the flash crash phase, despite the huge variation in prices.

⁹In the computation of the mean of the Poisson distribution, the relevant market volumes are weighted by the parameter $0 < \lambda < 1$. Earlier works also used Poisson distribution to represent order placement and cancellation (Zovko and Farmer, 2002; Farmer et al, 2005; Paddrik et al, 2011). Furthermore, results are likely to be robust to other types of distribution characterizing order size (Mike and Farmer, 2008).

¹⁰Our assumption about HF orders' size reflects empirically-observed HF characteristics, namely HF traders are few firms in the market but represent more than 30% of total trading volume (Kirilenko et al, 2011; Aldridge, 2013).

¹¹This assumption is consistent with empirical evidence on HF agents' behavior, which suggests that

Current market conditions											
Bid						Ask					
8	8	9	9	10		11	11	11	12	12	
HF trading											
Buy orders											
Bid						Ask					
8	8	9	9	10	12	11	11	11	12	12	
Sell orders											
Bid						Ask					
8	8	9	9	10		9	11	11	11	12	12

Figure 1: Directional strategy order placement.

are formed as follows:

$$P_{j,t} = P_{j,t}^{ask}(1 + \kappa_j) \quad P_{j,t} = P_{j,t}^{bid}(1 - \kappa_j), \quad (8)$$

where κ_j is drawn from a uniform distribution with support $(\kappa_{min}, \kappa_{max})$. Figure 1 illustrates HF order placement using directional strategy based on a spread of one tick. Notice that, although the above limit pricing strategy is meant to minimize non-execution risk, the structure of the HF limit order placement in the model implies that HF orders' execution is not always guaranteed, not even partial one. Indeed, as we explained at the beginning of Section 2 HF order submission takes place sequentially, and the best bid and the best ask can change as HF traders consecutively place their orders in the book. It follows that the limit price (*e.g.*, a bid quote) set by one given HF trader can be outperformed by the one set by another HF trader in the sequence, thereby partially or totally compromising the execution of the first order.¹²

A key characteristic of empirically-observed high-frequency trading is the high order cancellation rate (CFTC and SEC, 2010; Kirilenko et al, 2011). We introduce such a feature in the model as follows. In each period, HF traders are able to process and to use available market information to decide whether to cancel their orders, while these

most of their orders are placed very close to the last best prices (SEC, 2010).

¹²In addition, the probability that a given HF order is executed in the final matching procedure is further affected by the possibility of order's cancellation (see also below in the section). The fact that HF orders are not systematically executed at the best available price is also confirmed by the analysis of agents' order aggressiveness. Table 3 clearly indicates that during normal times HF aggressiveness ratios are low and comparable to the ones of LF agents (or lower). The same occurs during flash crash recoveries. It follows that, most of the time, HF limit orders do provide liquidity to the market and are not aggressive orders.

orders could be executed. More precisely, we assume that, once new orders have been inserted in the book, HF traders are able to simultaneously compute the volumes and prices of the transactions that would take place conditionally on the existing book, and on this basis, they are able to compute their expected profits. HF agents will cancel their orders, and hence will choose not to participate in the market in that period, when expected profits are negative. Instead, they will confirm their orders, and will decide to take part in the trading process, when expected profits are non-negative. More formally, let $\pi_{j,t}^E$ be the expected profits of HF trader j conditional on the book available in period t , we get:

$$\begin{cases} \pi_{j,t}^E < 0, & \text{cancel order} \\ \pi_{j,t}^E \geq 0, & \text{confirm order} \end{cases}$$

where $\pi_{j,t}^E$ is determined by:

$$\pi_{j,t}^E = (\bar{P}_t^{temp} - P_{j,t}^{temp})D_{j,t}. \quad (9)$$

where \bar{P}_t^{temp} and $P_{j,t}^{temp}$ are the temporary market price and the temporary transaction price of agent j , respectively, and $D_{j,t}$ is the size of the order of agent j .

Note that, although HF traders get activated (or not) on the basis of past price changes (see equation 7), they actually confirm their orders and trade based on the expected profits that could be reaped in that period, given the state of the book.¹³

Lastly, at the end of each trading session, HF traders' profits ($\pi_{j,t}$) are computed as follows:

$$\pi_{j,t} = (\bar{P}_t - P_{j,t})D_{j,t}. \quad (10)$$

where $P_{j,t}$ is her actual transaction price and \bar{P}_t is the actual market price. As already mentioned at the beginning of the section, given that HF traders can intentionally cancel orders, the final order book would be different from the one before HF traders' endogenous cancellation. Accordingly, expected profits of HF traders could be different from expected ones.

¹³This representation of HF traders' behavior enables us to account for the complexity of such a behavior and in particular for HF low-latency and for HF market-driven participation in the market (*i.e.*, depending both on price movements and observed profitable opportunities). This view of HF traders behavior has also received support in recent empirical studies (see, for instance, CFTC and SEC, 2010; Kirilenko et al, 2011; Easley et al, 2012).

3 Simulation Results

We investigate the properties of the model presented in the previous section via Monte-Carlo simulations. More precisely, we carry out $MC = 50$ Monte-Carlo iterations, each one composed of $T = 1,200$ trading sessions using the baseline parametrization, described in Table 1. The value of the parameters employed in our simulations are in line with existing works.¹⁴

Table 1: Parameters values in the baseline scenario

Description	Symbol	Value
Monte Carlo replications	MC	50
Number of trading sessions	T	1,200
Number of low-frequency traders	N_L	10,000
Number of high-frequency traders	N_H	100
LF traders' trading frequency mean	θ	20
LF traders' min and max trading frequency	$[\theta_{min}; \theta_{max}]$	[10;40]
Chartists' order size parameter	α_c	0.04
Chartists' shock standard deviation	σ^c	0.05
Fundamentalists' order size parameter	α_f	0.04
Fundamentalists' shock standard deviation	σ^f	0.01
Fundamental value shock standard deviation	σ^y	0.01
Fundamental value price drift parameter	δ	0.0001
LF traders' price tick standard deviation	σ_z^L	0.01
LF traders' intensity of switching	ζ	1
LF traders' resting order periods	γ^L	20
HF traders' resting order periods	γ^H	[1;1,200]
HF traders' activation threshold distribution support	$[\eta_{min}; \eta_{max}]$	[0;0.2]
Market volumes weight in HF traders' order size distribution	λ	0.625
HF traders' order price distribution support	$[\kappa_{min}; \kappa_{max}]$	[0;0.01]

Figures 2 and 3 provide an example of, respectively, the evolution of price and price returns generated by our model. Both, the one in Figure 3 in particular, indicate the presence of extreme fluctuations in the simulated data, as well as the presence of volatility clusters. As a first step in our analysis of simulation results, we verify that our ABM is able to jointly reproduce the main stylized facts of financial markets, also detected at very high frequency (*i.e.*, less than one day, see Section 3.1). We then assess the properties

¹⁴More precisely, for the LF trading strategies equations, we chose the same values employed in previous ABM works (*e.g.*, Westerhoff, 2008). In addition, following Paddrik et al (2011), several values of the parameters concerning HF traders' behavior (*e.g.*, order size) were selected in order to be consistent with the evidence reported in Kirilenko et al (2011).




Figure 2: Evolution of the price in a Monte-Carlo simulation run. Time scale: minutes. y-axis in logarithmic scale.




Figure 3: Evolution of 1-minute returns in a Monte-Carlo simulation run.

of the model in generating flash crashes and we investigate the key determinants of flash crashes, distinguishing the initial sharp price drop and the subsequent price recovery (see Section 3.2). Lastly, we investigate the effectiveness of a set of regulatory policies on market volatility, the frequency and the duration of flash crashes (see Section 3.3).

3.1 Stylized Facts of Financial Markets

We follow an indirect calibration approach to the validation of our agent-based model (see Windrum et al, 2007, for a discussion of this approach) by checking its ability to jointly reproduce several stylized facts of financial markets with the same configuration of parameter values.

First, in line with empirical evidence (*e.g.*, Cont, 1997; Bollerslev and Wright, 2000; Mills, 2001; Chakraborti et al, 2011), we find that our model generates zero autocorrelation values of price-returns (calculated as logarithmic differences, see Figure 4).¹⁵ In contrast to price returns, the autocorrelation functions of absolute returns display a slow decaying pattern (cf. Figure 5), thus confirming the presence of volatility clustering in our simulated data (Cont et al, 1997; Cont, 2001; Mills, 2001; Chakraborti et al, 2011).

Another widely-studied property of financial markets is the presence of fat tails in the distribution of price returns (Abhyankar et al, 1995; Mantegna et al, 1995; Gopikrishnan et al, 1999; Mills, 2001; Gorski et al, 2002; Chakraborti et al, 2011). We plot in Figure 6 the density of pooled returns across Monte-Carlo runs (stars) together with a normal

¹⁵Note that consistent with our finding, empirical evidence suggest that autocorrelations of returns are mainly insignificant, except for very small intraday time scales (*e.g.*, Cont et al, 1997; Cont, 2001; Mills, 2001; Selçuk and Gençay, 2006), but rapidly decay to zero in few minutes so that they can be safely assumed to be zero (Cont et al, 1997). Earlier empirical works explain this phenomenon through microstructure effects (see, for instance, Cont, 2001; Selçuk and Gençay, 2006).




Figure 4: 1-minute price-returns sample autocorrelation function (solid line) together with 95% confidence bands (dashed lines). Values are averages across 50 independent Monte-Carlo runs.




Figure 5: Sample autocorrelation functions of absolute 1-minute price returns (solid line) together 95% confidence bands (dashed lines). Densities are estimated using a kernel density estimator using a bandwidth optimized for Normal distributions.




Figure 6: Density of pooled 1-minute price returns (stars) across 50 independent Monte-Carlo runs together with a Normal fit (solid line). Logarithmic scale on y-axis.




Figure 7: Sample autocorrelation functions of absolute 1-minute price returns (solid line) together 95% confidence bands (dashed lines). Densities are estimated using a kernel density estimator using a bandwidth optimized for Normal distributions.

density (solid line) fitted on the pooled sample. As the figure shows, the distribution of price returns significantly departs from the Gaussian benchmark. Moreover, Figure 7 shows the tail of the distribution of (negative) price returns together with a power-law fit.¹⁶ Again, and in line with empirical evidence (Mantegna et al, 1995; Gopikrishnan et al, 2000; Gorski et al, 2002), the power law distribution provides a good approximation of the simulated data of tail returns.




Figure 8: Evolution of the price during a flash crash. Time scale: 1 minute. Values of the price on the y-axis are expressed in relation to the price level in the period just before the crash.

3.2 HFT and the Anatomy of Flash Crashes

In line with empirical evidence (CFTC and SEC, 2010; Kirilenko et al, 2011), we identify flash crashes as drops in the asset price of at least 5% followed by a sudden recovery of at most 30 minutes (corresponding to thirty trading sessions in each simulation run). Applying such a definition, in line with Jacob Leal et al (2016), we find that our model is able to endogenously generate flash crashes as an emergent property resulting from the interactions between low- and high-frequency traders. Figure 8 provides a visual example of flash crash in our simulations, by showing the price evolution just before and then after the crash. Moreover, we find that flash crashes emerge only when HF traders are present in the market and their frequency is significantly higher than one (see Table 2). In contrast, when the market is only populated by LF traders, flash crashes do not

¹⁶The power-law exponent was estimated using the freely available “power-law package” and based on the procedure developed in Clauset et al (2009).

Table 2: Market volatility (σ_P) and flash crashes statistics in the baseline scenario with HF traders and in the scenario with only LF traders.

	σ_P	Number of flash crashes	Avg. duration of flash crashes
Baseline	0.016 (0.002)	4.636 (0.398)	7.139 (0.484)
Only-LFT	0.002 (0.000)	- -	- -




Figure 9: Complementary cumulative distributions of bid-ask spreads in different market phases. Pooled sample from 50 independent Monte-Carlo runs.

emerge.

What are the main drivers of the emergence of flash crashes in our model?

First, the directional strategies employed by HF traders can lead to large bid-ask spreads, setting the premises for the emergence of flash crashes. Figure 9 shows the distributions of bid-ask spreads conditioned on different market phases *i.e.*, normal times, crash and recovery phases.¹⁷ We observe that the mass of the distribution of bid-ask spreads is significantly shifted to the right during crashes, clearly indicating the presence of large

¹⁷In particular, we construct the pooled samples (across Monte-Carlo runs) of bid-ask spread values singling out “normal time” phases and decomposing “flash-crash” periods in “crash” phases (*i.e.* periods of sharp drops in the asset price) and the subsequent “recovery” phases (*i.e.* periods when the price goes back to its pre-crisis level). Next, we estimate the complementary cumulative distributions of bid-ask spreads in each market phase using a kernel-density estimator.

bid-ask spreads at the time of the price fall. The emergence of large bid-ask spreads is explained by the different strategies employed by high- and low-frequency traders in our model. Active LF traders set their order prices “around” the price of the last trading session, which tends to fill the existing gap between the best bid and ask prices at the beginning of a given trading section. In contrast, active HF traders, who submit their orders after LF agents, place large buy (*sell*) orders just few ticks above (*below*) the best ask (*bid*), which tends to occasionally generate large bid-ask spreads in the LOB.¹⁸

HFT-induced large bid-ask spreads are therefore one key driver of flash crashes in our model. This is further confirmed by the analysis of agents’ aggressiveness behavior conditional on different market phases. First, Figure 10 shows that the model generates a positive relationship between agents’ orders aggressiveness and the size of the bid-ask spreads.¹⁹ Thus, higher orders’ aggressiveness (of any agents’ type) leads to larger bid-ask spreads in our model. However, the degree of aggressiveness of HF and LF agents markedly differ across market phases. Table 3 shows average orders’ aggressiveness ratios²⁰ for both types of agents and book sides, and conditional on different market phases (*i.e.*, “normal times”, “crash”, “recovery”). As the table shows, in all market phases, aggressiveness ratios of LF traders are low in all market phases. Moreover, order aggressiveness of HF traders is low during normal times: respectively 87% of buy and 89% of sell orders placed by HF traders do provide liquidity to the market, which contributes to keep bid-ask spreads low. In contrast, orders’ aggressiveness of HF agents increases abruptly in the crash phase (see Table 3). In particular, in such a phase, most HF sell orders are aggressive (about 85%) and thus remove liquidity from the market and generate large bid-ask spreads. Instead, HF aggressiveness is very low on the buy side.²¹

¹⁸Moreover, in the appendix we also report results of experiments where we assume that the pricing strategy of HF traders is the same as the one of LF traders. These results indicate that high volatility and flash crashes emerge also in that case, as long as the standard deviation of the price-tick distribution of HF traders is enough larger than the one of LF traders.

¹⁹Figure 10 shows the contour of the theoretical function between bid-ask spreads and HF and LF orders’ aggressiveness ratios that is implied by the dynamics of the model. The function was generated by interpolating the scattered data of aggressiveness ratios of HF and LF traders and bid-ask spread pooled across Monte-Carlo simulations in the baseline scenario. The interpolation was performed by using the *scatteredInterpolant* function in Matlab.

²⁰Alike Jacob Leal et al (2016), we use the definition provided by trading platforms (*e.g.*, CME Globex), and widely used in the empirical literature (Kirilenko et al, 2011; Baron et al, 2014). An incoming order is considered “aggressive” if it is matched against an order that is resting in the book, *i.e.*, if it removes liquidity from the market. In contrast, an order provides liquidity on the market if it fills the book of resting orders. Finally, it has no effect on market liquidity if it is matched against another incoming order in the same trading session.

²¹In Jacob Leal et al (2016), we also show that such an asymmetry is further confirmed by the distribution of overall orders of HF traders across market sides. There, we also explain how HF orders’




Figure 10: Contour plot of the relation between agents' order aggressiveness and size of the bid-ask spread generated by the model.

Table 3: Orders' aggressiveness ratios for different categories of traders and different market phases. Values are averages across 50 independent Monte-carlo runs. Monte-carlo standard errors in parentheses.

	LFT buy	HFT buy	LFT sell	HFT sell
Normal times	0.086 (0.003)	0.130 (0.002)	0.041 (0.003)	0.108 (0.002)
Crashes	0.000 (0.000)	0.013 (0.008)	0.000 (0.000)	0.831 (0.031)
Recovery	0.083 (0.013)	0.106 (0.009)	0.004 (0.002)	0.095 (0.010)
Unconditional values	0.086 (0.003)	0.130 (0.002)	0.041 (0.003)	0.110 (0.002)

To sum up, the above discussion shows that flash crashes in our model are the result of: *i.*) large bid-ask spreads occasionally generated by HF traders’ pricing strategies and *ii.*) concentration of aggressive orders on the sell-side of the book. It is worth noticing that these explanations for the emergence of flash crashes are in line with the empirical evidence about the market dynamics observed, for instance, during the flash crash of May 6th, 2010 (CFTC and SEC, 2010; Kirilenko et al, 2011). Moreover, computer simulations highlight the key role that high-frequency trading has in generating such extreme events in financial markets. Indeed, the emergence of periods of high market illiquidity is endogenous and intimately related to the pricing strategies of HF traders (see Eq. 8). In that, flash crashes are therefore not simply generated by large orders and thus cannot be associated with “fat finger” explanations (see Haldane, 2014, for a discussion of the different proposed explanations of flash crashes). Finally, our simulation results confirm on the one hand that, in line with recent empirical evidence (see *e.g.*, Brogaard, 2010; Menkveld, 2013), HF traders may have a beneficial effect on markets during normal times, by providing non-aggressive orders and therefore contributing to keep bid-ask spread low. On the other hand, they also show that the liquidity provided by HF traders is extremely fragile and that their orders can occasionally be extremely aggressive, removing liquidity from the market, and generate abrupt and large drops in the market price.²²

The above discussion has made clear that HFT plays a key role in causing the significant price falls, which are the footprint of all flash crashes. However, HFT also actively contributes to the quick recovery after the crash. Table 3 shows indeed that the orders’ aggressiveness ratios of HF agents are much lower during the recovery phase of the flash crash. In addition, orders’ aggressiveness ratios are symmetric between the buy and sell sides of the book. Thus, orders of HF agents contribute to restore liquidity in such a phase, thereby favoring the recovery of the price. The return to normal liquidity conditions during the recovery is also documented by the behavior of the conditional

synchronization is an emergent property related to the event-time strategy of HF traders and may emerge even if the choice of each HF agent between selling or buying is a Bernoulli distributed variable with probability $p = 0.5$.

²²Note that, in order to check the robustness of our results, we also ran additional simulation exercises where we changed the value of the parameter η_{max} , *i.e.*, the upper bound of the support of the uniform distributions from which agents’ idiosyncratic activation thresholds are drawn, from 0 to 0.4. In the baseline scenario, instead the value of this parameter is $\eta_{max} = 0.2$, which implies significant variability across agents in terms of activation thresholds, as well as activation based on both small and large price changes. We find that reducing the η_{max} with respect to the baseline has no significant effect on volatility. but , it significantly increases the number of flash crashes (and their duration). Also increasing the variability of thresholds increases volatility and flash crashes, although to a much lesser extent. These results are available from the authors upon request.

bid-ask spread distribution (cf. Figure 9). The distribution of the bid-ask spreads during recoveries is indeed not statistically different from the one during normal times.

Two factors explain the positive role played by HFT in favoring price-recovery after the crash. The first is that wide variations in asset prices trigger the activation of a large number of high-frequency traders which leads to a surge in order volumes of HF agents.²³ In addition, as each HF trader is either a buyer or a seller with probability $p = 0.5$, when the number of active HF agents is large, HF orders will tend to be equally split between the sell- and buy-side of the LOB, which explains the symmetry in HF orders' aggressiveness ratios observed during the recovery (see Table 3). The second element supporting the rapid price recovery is the order-cancellation rate of HF traders. Indeed, high order cancellation implies a short duration of HF orders in the book. As a matter of fact, this also implies that the HF bid and ask quotes will tend to reflect current market conditions. Such a *memory effect* of HF orders explains the low time persistence of high bid-ask spreads after a crash and contributes to the quick replenishment of market liquidity and price.

3.3 Regulatory policies experiments

In the previous section, we have documented that the model is able to robustly reproduce the main stylized of financial markets and to endogenously generate flash crashes as the result of the trading activity of HF traders. We pointed out the very reasons underlying both phases of a flash crash, namely the sharp price drop and the swift recovery of the price. In this context, we now turn to use the model as a test-bed to investigate the effectiveness of a set of regulatory policies which have been so far implemented and proposed to cope with the possible negative effects of high-frequency trading and to curb flash crashes. We focus on the following policies: *i.*) circuit breakers, *ii.*) minimum resting times, *iii.*) cancellation fees, *iv.*) financial transaction taxes. Moreover, we study the impact of the aforementioned policies on price-returns volatility as well as on the number and duration of flash crashes. Our focus on these policy goals is motivated first by the fact that they received much attention in recent debates (see e.g. Haldane, 2014). Moreover, in Jacob Leal et al (2016), we developed a model explaining how HFT strategies could exacerbate volatility and generate flash crashes, and we also identified the factors affecting their incidence and duration. It is then a natural extension to analyze how regulatory policies could affect the foregoing aspects of market stability.

²³In the appendix we also report results about experiments for the case where agents activate only in presence of small price changes and we show that the main results of our model robustly hold.

3.3.1 Circuit breakers

Section 3.2 provides insights about the mechanisms through which HFT may be a source of episodic price instability and systemic risk. Regulators have recently taken proactive steps to avoid flash crashes and to deal with periodic illiquidity in markets. In particular, in the aftermath of May 2010 Flash Crash, the CFTC and the SEC proposed several measures to preclude this type of extreme events such as, for instance, updated circuit breakers (SEC, 2011b, 2012) and limit up/limit down mechanisms (also known as ex-ante circuit breakers, see SEC, 2011a, 2012; Haldane, 2014). Indeed, extreme price fluctuations are likely to exacerbate execution uncertainty and discourage trading (Greenwald and Stein, 1991; Subrahmanyam, 2012). Instead, trading halts should allow for a “cool-down” period, improve market liquidity and reduce volatility (Greenwald and Stein, 1991; Kodres and O’Brien, 1994; Ackert, 2012). Circuit breakers (or impediments to trade), *i.e.*, mechanisms designed to reduce the risk of a price collapse by means of trading halts in presence of excessive price volatility, have been implemented for long time in many exchanges, both in Europe and in the US (CFTC and SEC, 2010; Furse et al, 2011; Prewitt, 2012; Gomber and Haferkorn, 2013). However, they were traditionally market-wide and triggered only by large price movements. They were therefore conceived only as ex-post reactions to excessive price volatility. After the events of May 6, 2010, new and more sensitive stock-specific systems, which work on an ex-ante basis, have therefore been implemented (SEC, 2012). Nowadays, circuit breakers can take many forms, from trading halts in single stocks or in entire markets to limit up and down prices with a variety of percent price change and different reference points, and restrictions on one trading venue or across multiple venues (Furse et al, 2011; Subrahmanyam, 2012).

However, on the one hand, the empirical evidence on the efficacy of circuit breakers²⁴ and price limits²⁵ is so far limited. On the other hand, from a theoretical viewpoint, the debate about the effectiveness of trading halts reveals that these devices may have either positive (Westerhoff, 2003b, 2006, 2008) or/and negative effects (Yeh and Yang, 2010, 2013). Furthermore, this debate has mainly focused on the levels of price limits rather than the effectiveness of different types of circuit breakers. Consequently, it is not clear yet what type of breakers are the most effective to deal with flash crashes. In this section, we try to contribute to the existing literature on circuit breakers by performing

²⁴Few examples include *e.g.*, Lauterbach and Ben-Zion (1993); Santoni and Liu (1993); Goldstein and Kavajecz (2004); Brugler and Linton (2014).

²⁵Some examples include, for instance, Kim and Rhee (1997); Cho et al (2003); Diacogiannis et al (2005); Bildik and Gülay (2006); Stamatiou (2008).

a computational test of their impact on volatility and the duration of flash crashes. We




Figure 11: Evolution of the price during a flash crash in a market with ex-post trading halts (solid line) and in the unregulated market (dashed line). Values of the price on the y-axis are expressed in relation to the price level in the period just before the crash.

focus on two distinct types of circuit breakers,²⁶ namely: *i.* a trading halt in a single stock triggered by a given percent price change from the last price, as implemented, for instance, at the NYSE-Euronext, the London Stock Exchange and the Deutsche Bourse (*i.e.*, an ex-post device); *ii.* a limit up/limit down price mechanism in place *e.g.*, on the NYSE and NASDAQ, Tokyo Stock Exchange and Korea Exchange (*i.e.*, an ex-ante device). More precisely, we first study the effect of introducing an ex-post trading halt mechanism in response to substantial price drops which is intended to stop trading in the exchange for a time period (np). In this Monte-Carlo experiment, the circuit breaker is triggered by a relative price change from the last price that is in absolute value equal or larger than $\beta\%$, where $\beta = 5\%$,²⁷ and the trading halt is assumed to last for $np = 5$ periods. Notice that, by construction, this type of circuit breaker leaves unaffected the number of flash crashes, but it could have an impact on the duration of flash crashes and on market volatility. Figure 11 (solid line) provides a visual example of the dynamics of flash crash in presence of ex-post trading halts in our simulation experiments. The same

²⁶Note that in our model structure does not assume any specific information disclosure about the type of circuit breaker that is implemented in the market, which allows us to discard the strategic reaction of traders and to only focus on the effect of discretionary circuit breakers, when agents hardly trade in anticipation of a market halt.

²⁷Notice that, in our parametrization, the threshold for the trading halt activation corresponds to the one used to identify flash crashes.

Table 4: The effect of different types of circuit breakers on price volatility and flash crash statistics when $\gamma^H = 20$ and $\beta = 5\%$. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

	σ_P	Avg. duration of flash crashes
No circuit breaker baseline	0.016 (0.002)	7.139 (0.484)
ex-post circuit breaker	0.010 (0.001)	13.345 (0.609)
ex-ante circuit breaker	0.005 (0.000)	- -

figure also compares the flash crash dynamics with trading halts to the one emerging in the unregulated market. In the second type of experiment, we introduce a limit up/limit down price of $\beta = 5\%$ which, when triggered, stops trading in the exchange for $np = 5$ periods. In this case, the trading halt occurs before the trading session price is formed. In such a case, the imposition of limit up/limit down price completely removes flash crashes from the dynamics.

The results of both experiments are shown in Table 4.²⁸ We only report results for market volatility and the average duration of flash crashes. First, we find that the introduction of ex-post circuit breakers has a negligible effect on market volatility. Moreover, this defensive regulation has a detrimental effect on the duration of the flash crash, since the trading halt merely slows down the price recovery. Indeed, we observe that the price would have recovered sooner without the imposition of the circuit breaker, *i.e.*, if HF traders would have been able to fully play their role in the recovery phase of the flash crash. This is mainly explained by the positive role played by HF traders in the recovery from the crash (see Section 3.2). The imposition of a trading halt instead prevents HF traders from providing the required liquidity after the crash and thus leads to longer flash crashes.

How do results change if we turn from ex-post to ex-ante circuit breakers? Besides removing flash crashes, in our parametrization, in line with earlier works (*e.g.*, Westerhoff, 2008; Yeh and Yang, 2013), trading halts lead to a reduction in price volatility

²⁸We also ran simulations for alternative values of γ^H ($\gamma^H = 1$ and $\gamma^H = 1,200$) and for $\beta = 3\%$. Results are however consistent with the ones presented in Table 4.

compared to the baseline (compare first and third row of the first column in Table 4). In our setting, this is mainly explained by the fact that this device is triggered before trade is actually performed and therefore it prevents extreme price fluctuations and HFT activity. As a result, trading halts that hinder the market from collapsing are relevant and effective ways to deal with flash crashes and to curb HFT.

Overall, and in line with earlier works (see *e.g.*, Subrahmanyam, 2013; Yeh and Yang, 2010, 2013; Apergis, 2014), our results show that breakers should be used with caution, especially when they represent impediments to trade and deteriorate the trading process within a particular stock. In particular, our findings indicate that ex-ante circuit breakers are a much more effective tool than ex-post trading halts, because, consistent with Westerhoff (*e.g.*, 2003b, 2006, 2008), they hamper extreme drops in price from the market and they significantly dampen market volatility. In contrast, ex-post trading halts have only a limited impact on volatility. In addition, as suggested by Fama (*e.g.*, 1989), they may introduce important distortions in the natural process of recovery from a crash that would otherwise took place and merely delay the price discovery.

3.3.2 Minimum resting times

Minimum resting times specify a minimum time that a limit order must remain in the book *i.e.*, it cannot be cancelled within a given time span. The impetus for imposing this command-and-control regulatory instrument is that markets operating at high speed are characterized by a large number of orders that are cancelled very quickly after submission. Orders' cancellation is an inherent feature of many HF traders strategies and has raised many critiques against HFT (CFTC and SEC, 2010; Kirilenko et al, 2011). Indeed, the ability of HF traders to quickly cancel their orders could render market liquidity misleading (Kirilenko et al, 2011; Prewitt, 2012; Breckenfelder, 2013; Friederich and Payne, 2015), and it could favor price short-term volatility (Hanson, 2011; Bershova and Rakhlin, 2013; Breckenfelder, 2013). Furthermore, rapid order cancellations are likely to increase the cost of monitoring the market for all participants and reduce the predictability of a trade's execution quality, given that the quotes displayed may have been cancelled by the time the new order hits the resting order (Furse et al, 2011). Nevertheless, the net benefits of minimum resting times are still unclear (Furse et al, 2011). On the one hand, minimum resting times can increase the likelihood of a viewed quote being available to trade and therefore make the order book dynamics more transparent. In addition, longer expiration times create liquidity that reduces price variance in the market (Brewer et al, 2013). Lastly, by "slowing down" markets, minimum resting

times may favor participation, especially if some traders (*e.g.*, small retail investors) feel that high speed makes market unfair and hurts market integrity (see, for instance, Haldane, 2014). On the other hand, minimum resting times can impinge upon hedging strategies that operate by placing order across markets and expose liquidity providers to increased “pick-off risk” due to the inability to cancel stale orders (for Science, 2012). Liquidity provision may be even more impeded during times of high volatility, when it is particularly expensive to post limit orders. Furthermore, this measure may also change the dynamics of the market by attracting more aggressive HFT (Farmer and Skouras, 2013). Lastly, market quality may be diminished due to higher transaction costs for the end users and lower price efficiency.




Figure 12: Market volatility as a function of minimum resting times duration. Values on the y-axis are averages across 50 independent Monte-carlo runs and are expressed in relation to the baseline value in the unregulated market.




Figure 13: Average number (solid line, left scale) and average duration of flash crash (dashed line, right scale) as a function of minimum resting times duration. Values on the y-axis are averages across 50 independent Monte-carlo runs and are expressed in relation to the baseline value in the unregulated market.

In this context, and given that the empirical evidence about the effects of minimum resting times is still limited,²⁹ we aim at shedding some light on the impact of minimum resting times on market dynamics by investigating the effects of such a measure on market volatility as well as on the number and the duration of flash crashes. To this end, we run a Monte-Carlo experiment where we impose that HF orders cannot deliberately cancel their orders for a number of periods equal to the expiration time γ^H .³⁰ We then vary the

²⁹For instance, the work of Furse et al (2011) reports only two cases for the implementation of minimum resting times. Namely ICAP which introduced a minimum quote lifespan on its electronic in June 2009 and the Istanbul Stock Exchange which did not allow the cancellation of limit orders during continuous auction mode until mid-2011. However, it is not clear what one can really learn from these two experiments.

³⁰Given that HF agents’ unexecuted orders are automatically removed from the book after a given time period γ^H , in this experiment, it is sufficient to neutralize the HF endogenous cancellation procedure

Table 5: HF traders' minimum resting times, price volatility and flash crash statistics. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

γ^H	σ_P	Number of flash crashes	Avg. duration of flash crashes
1	0.017 (0.002)	7.114 (0.845)	9.527 (0.746)
3	0.005 (0.000)	1.556 (0.103)	10.537 (0.834)
5	0.005 (0.000)	1.143 (0.053)	13.929 (1.476)
10	0.004 (0.000)	1.250 (0.071)	13.500 (1.577)
15	0.003 (0.000)	1.000 (0.000)	13.000 (1.497)
20	0.007 (0.001)	1.000 (0.000)	26.333 (0.898)
40	0.002 (0.000)	- -	- -
60	0.002 (0.000)	- -	- -

expiration time while keeping all other parameters at their baseline values (see Table 1). In this experiment, we change the parameter γ^H from 1 to 60 periods/minutes. The results of this experiment are reported in Table 5. Table 5 reveals that increasing minimum resting times (*i.e.*, making γ^H higher) dampens market volatility (see second column in Table 5). The beneficial effect on volatility is one of the purported primary effects of the measure (see in particular SEC, 2010) and is consistent with earlier works (see in particular Hayes et al, 2012). This outcome is mainly explained by the fact that minimum resting times slow down HF traders and prevent them from aggressively trading on the most recent news and information disclosed in the LOB.

Minimum resting times (*i.e.*, higher γ^H) have also a beneficial effect on the number of flash-crash episodes (see third column of Table 5). This outcome again stems from and to vary the value of the parameter γ^H to test the effectiveness of minimum resting times.

the lower aggressiveness that such a measure imposes on HF trading strategies. In contrast, we find that the duration of flash crashes is inversely related to the duration of minimum resting times (cf. fourth column of Table 5). Figures 12 and 13 summarize the main findings discussed so far, by showing how price volatility, on the one hand, and the number and duration of flash crashes on the other hand vary with the length of minimum resting times. In particular Figure 13 provides a visual idea of the trade-off between flash crash incidence and its duration.

The longer flash crash duration observed for higher values minimum resting times is explained by the fact that stricter rules on orders' expiration of HF orders also imply a longer memory effect (cf. Section 3.2). In fact, as γ^H increases, the bid and ask quotes posted by HF agents stay longer in the LOB and therefore large bid-ask spreads persist more. Furthermore, less HF traders participate in the market based on the most recent market information. This slows down the replenishment of market liquidity and prevents the quick price recovery. Lastly, the number of contracts traded at prices close to the flash-crash one rises which prevents the price rebound.

Overall, the above results thus indicate that the imposition of minimum resting times can be a very effective tool in order to dampen market volatility and to reduce the incidence of flash crashes. In that, they bring support to earlier works advocating for such a measure (Haldane, 2014; SEC, 2010). At the same time, they also hint to the presence of a trade-off between volatility and incidence of extreme events, on the one hand, and price-resilience (because of longer recoveries) on the other hand.³¹ As we shall discuss in the next sections, such a trade-off is also inherent the market-based measures on which we focus on, namely cancellation fees and financial transaction taxes.

3.3.3 Cancellation fees

We now turn to investigate the effect of the imposition of cancellation fees on price volatility, the frequency of flash crashes and their duration. Both US and EU regulators have called for the imposition of cancellation fees. However, they have been only incompletely enforced in a couple of exchanges since 2012 (Nasdaq and Direct Edge, Borsa Italia and Deutsche Börse stock exchanges). Cancellation fees are primarily intended to prevent overload in the exchange computer systems and to discourage the most flagrant excessive cancellations which represent unnecessary messages that do not result in trades and which, rather, come along with higher volatility (Prewitt, 2012). A portion of such traffic is likely to be inefficient and may raise costs to other investors who try to mon-

³¹Haldane (2014) also points to the presence of a similar trade-off when deciding whether to impose resting rules or not (market efficiency versus stability).

itor the market. Such fees would therefore discourage traders from posting orders that are not intended to be executed (Prewitt, 2012). They will also discourage manipulative HFT strategies (like stuffing and spoofing) that involve massive order cancellations by rendering them uneconomical (Biais and Woolley, 2011; Prewitt, 2012). At the same time, rapid reaction to new information is often a way for market makers to minimize the risks of offering prices to other traders, and contributes to lower trading costs (Copeland and Galai, 1983; Foucault et al, 2003). In that, the imposition of cancellation fees could instead discourage the activity of active market makers and liquidity providers, and lead to an increase in transaction costs.




Figure 14: Market volatility for different levels of cancellation fees and expiration dates γ^H . Bars heights correspond to average Monte-Carlo values in relation to the baseline value in the unregulated market under each configuration of γ^H .

In our experiment, HF traders who deliberately decide to cancel their orders before γ^H periods are charged a fee c . As a result, a HF agent will cancel her order if expected losses from trade are higher in magnitude than the cancellation cost, *i.e.*, when $\pi_{j,t}^E < -c.D_{j,t}$. The policy exercise was carried under three scenarios: *i.*) when $\gamma^H = 1$, HF traders can deliberately decide to cancel their orders before the expiration date (γ^H periods). However, given that the expiration date, in this case, is very small (*i.e.*, $\gamma^H = 1$), HF traders will have to pay the cancellation fee only on very fast order cancellation. This scenario represents a very soft policy measure where most cancelled orders are not charged the fee; *ii.*) when $\gamma^H = 20$, HF traders can decide to deliberately cancel their orders before their expiration date (γ^H periods), which leads to the payment of the cancellation fee c . However, older unexecuted HF orders which are automatically withdrawn from the book after γ^H are not charged the cancellation fee. This scenario

represents a moderate policy measure where not all cancelled orders are charged the fee; *iii.*) when $\gamma^H = 1, 200$ (i.e. it corresponds to the length of a Monte-Carlo iteration in our setting), HF orders can only intentionally be cancelled by HF agents. In this case, HF unexecuted orders stay in the LOB until the end of the simulation (i.e., $T = \gamma^H = 1, 200$) and the cancellation fee is charged on all cancelled HF orders. This scenario represents the imposition of a very stringent policy measure. Furthermore, given the wide variety of fee levels currently used worldwide, we tested the effect of different levels of cancellation fees, c varying from 0.01% to 1%. The results of the above experiments are reported in Table 6, and summarized by the bar plots in Figures 14 to 16.




Figure 15: Average number of flash crashes as a function of minimum resting times duration. Bars heights correspond to average Monte-Carlo values in relation to the baseline value in the unregulated market under each configuration of γ^H .




Figure 16: Average duration of flash crashes (dashed line, right scale) as a function of minimum resting times duration. Bars heights correspond to average Monte-Carlo values in relation to the baseline value in the unregulated market under each configuration of γ^H .

First, and not surprisingly, we find that, when $\gamma^H = 1$, the imposition of a cancellation fee is not effective in dealing with volatility and flash crashes, whatever the size of the cancellation fee is. Indeed, in this case, price volatility, the frequency and the duration of flash crashes are not significantly different with respect to the baseline case. This is mainly explained by the fact that, when $\gamma^H = 1$, HF traders frequently cancel their orders because they are not penalized by the cancellation fee.

In contrast, we find that, in scenarios (*ii.*) and (*iii.*), the introduction of a cancellation fee can be an effective tool to deal with market volatility and the number of flash crashes. Furthermore, the level of the fee matters, since we observe that the higher is the cancellation fee, the greater are the effects on price volatility and the occurrence of flash crashes. In particular, under the most stringent scenario (i.e., $\gamma^H = 1, 200$), flash crashes completely vanish for high values of the cancellation fee and this regulatory instrument is thus very effective to deal with such extreme events. In the mild scenario

Table 6: HF traders' order cancellation fees, price volatility and flash crash statistics for different values of γ^H and different values of c . Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

$\gamma^H = 1$	c	σ_P	Number of flash crashes	Avg. duration of flash crashes
	0	0.017 (0.002)	4.652 (0.390)	7.552 (0.575)
	0.01%	0.015 (0.001)	6.262 (0.568)	11.061 (0.632)
	0.1%	0.017 (0.002)	6.808 (0.793)	10.209 (0.576)
	1%	0.016 (0.002)	7 (0.720)	9.071 (0.568)
$\gamma^H = 20$	c	σ_P	Number of flash crashes	Avg. duration of flash crashes
	0	0.016 (0.002)	4.636 (0.398)	7.139 (0.484)
	0.01%	0.006 (0.000)	2.200 (0.238)	17.108 (0.957)
	0.1%	0.006 (0.001)	1.750 (0.123)	9.264 (1.234)
	1%	0.007 (0.001)	2.115 (0.231)	12.115 (1.162)
$\gamma^H = 1,200$	c	σ_P	Number of flash crashes	Avg. duration of flash crashes
	0	0.014 (0.001)	3.909 (0.389)	7.424 (0.531)
	0.01%	0.002 (0.000)	1.000 (0.000)	27.000 (0.200)
	0.1%	0.003 (0.000)	1.000 (0.000)	17.750 (1.794)
	1%	0.002 (0.000)	— —	— —

Table 7: Orders' aggressiveness ratios for different categories of traders and different market phases when $\gamma^H = 20$ and $c = ftt = 0.01$. Values are averages across 50 independent Monte-carlo runs. Monte-carlo standard errors in parentheses.

	Cancellation fee		Financial transaction tax	
	LFT orders	HFT orders	LFT orders	HFT orders
Normal times	0.157 (0.008)	0.036 (0.003)	0.133 (0.006)	0.159 (0.005)
Crashes	0.000 (0.000)	0.999 (0.000)	0.000 (0.000)	0.866 (0.039)
Recovery	0.332 (0.021)	0.109 (0.015)	0.069 (0.007)	0.215 (0.032)
Unconditional values	0.159 (0.009)	0.038 (0.003)	0.131 (0.006)	0.161 (0.005)

(*i.e.*, $\gamma^H = 20$) this type of policy measure is still effective to curb HFT and to mitigate flash crashes. These findings confirm one common claim against HFT according to which HF high cancellation rates may destabilize markets (SEC, 2014). Accordingly, preventing HF traders from quickly cancelling their orders decreases market volatility and completely removes flash crashes from the market.

Furthermore, the introduction of a cancellation fee tends to significantly reduce HF orders' aggressiveness. Table 7 shows the (buy and sell) orders' aggressiveness ratios for both HF and LF traders in the mild scenario when $\gamma^H = 20$ and $c = 0.01$. Reported values are unconditional and for different market phases. The table also compares orders' aggressiveness ratios with a cancellation fee to the ones that emerge in presence of a financial transaction tax (see next section). As this table reveals, the introduction of a cancellation fee generates a situation where the aggressiveness of HF traders is significantly lower than the one of LF traders, both unconditionally as well as in the normal times and recovery phases. Not surprisingly, when $\gamma^H = 20$, the average bid-ask spread is significantly lower than in the baseline (1.022 versus 1.577). These outcomes are explained by the fact that the existence of the cancellation fee effectively discourages HF traders to frequently cancel their orders, since they have an incentive to keep orders with a lower expected profit.

However, and similarly to minimum resting times and circuit breakers, the beneficial effects of cancellation fees come at the cost of a longer duration of flash crashes. Again, this outcome is explained by the fact that in presence of a cancellation fee, HF orders

stay longer in the book. This does not only prevent the activation of HF traders in the recovery, it also implies that HF quotes in the book tend to reflect close-to-crash conditions. We therefore point out that preventing HF traders from quickly modifying and cancelling their orders slows down the price recovery, since HF orders do not reflect the most recent market conditions. As a result, the positive role HFT plays in the recovery from the crash is significantly dampened. This is further supported by the fact that, when $\gamma^H = 20$, the trading to book volume ratio is significantly lower than in the baseline (0.067 versus 1.147).³²

Overall, we suggest that HF traders' high cancellation rates are harmful for the market since such a behavior favors market volatility and the occurrence of flash crashes. The imposition of a cancellation fee is effective in reducing market volatility and to mitigate flash crashes. Nevertheless, given the positive influence of HF traders during the recovery phase, this type of regulatory policy may prevent HF traders from participating to the recovery process and it may lengthen the duration of flash crashes.

3.3.4 Financial transaction taxes

To conclude our investigation of regulatory measures, we investigate the effects of the introduction of a financial transaction tax (FTT). So far different schemes and levels of taxes have been implemented all over the world. Examples are the stamp duty in the UK, the French financial transaction tax on high-frequency trading and the pricing scheme introduced on NYSE Euronext. In this work, we assume that HF executed orders are charged a fee $ftt > 0$. Accordingly, HF traders will intentionally cancel their orders whenever $\pi_{j,t}^E < ftt \cdot D_{j,t}$.

Although its recent introduction in some markets has mainly been motivated by the goal of raising revenues in response to major financial crises (staff, 2010; Pollin et al, 2003), financial transaction taxes have traditionally been indicated as a possible effective tool to discourage short-term speculation (Tobin, 1978), to curb negative effects of HFT practices and to improve systemic resilience of financial markets (Griffith-Jones and Persaud, 2012). Nevertheless, the effectiveness of a financial transaction tax is still a controversial and highly debated topic among academics (see, for instance, McCulloch and Pacillo, 2011, for a review of existing works on financial transaction taxes). On the one hand, empirical evidence on the relationship between FTT and market quality delivers mixed results (see, for instance, Roll, 1989; Umlauf, 1993; Jones and Seguin, 1997; Habermeier and Kirilenko, 2001; Hau, 2006; Gomber et al, 2015), although some

³²Note that values are averages across 50 independent Monte-Carlo runs.

studies (e.g. Colliard and Hoffmann, 2013) find that an FTT may have a permanent positive effect on low-latency trading, due to lower order aggressiveness and fewer rapid cancellations. On the other hand, many theoretical works suggest that an FTT can have a stabilizing effect (Ehrenstein, 2002; Westerhoff, 2003a, 2004; Westerhoff and Dieci, 2006).³³ However, other theoretical works also point out that such a stabilizing role is highly dependent on some important conditions such as market liquidity (Haberer, 2004), the level of the tax (Giardina and Bouchaud, 2004; Dupont and Lee, 2007; Demary, 2010; Fricke and Lux, 2015), the structure of the market (Pellizzari and Westerhoff, 2009). Lastly, many scholars view HFT as the main providers of liquidity in modern markets (Hendershott et al, 2011; Menkveld, 2013). In this view, a financial transaction tax would not be beneficial because it would hurt the functioning of markets and reduce market quality (Dupont and Lee, 2007).

We therefore contribute to the above debate by running Monte-Carlo experiments where we impose different levels of financial transaction tax as a percentage of HF orders' size, and we then investigate the resulting impact on market volatility as well as on the occurrence and the duration of flash crashes. Table 8 shows the results of this experiment when $\gamma^H = 20$.³⁴ This table shows that the introduction of an FTT has a beneficial impact on market stability and on the occurrence of flash crashes. When the FTT is implemented in the market, we observe a reduction in price volatility and in the number of flash crashes. Again, these positive effects come at the cost of a longer duration of flash crashes. Figures 17 and 18 help to visualize this last result by showing the evolution of, respectively, market volatility and of the number and duration of flash crashes as a function of different transaction tax levels.

Furthermore, the effectiveness of financial transaction taxes is much milder compared to other policy measures discussed so far (e.g. minimum resting times and cancellation fees). In particular, the reductions in volatility and in the number of flash crashes with respect to the baseline are much lower than the one obtained with cancellation fees of the same level as the tax (compare results in Table 8 to the results in Table 7 with the scenario $\gamma^H = 20$). Significant improvements are obtained only with draconian tax rates (*i.e.*, 10% or 50%, see Table 8). Moreover, the introduction of a financial transaction tax does not lead to lower HF orders' aggressiveness, as it was the case for the introduction of a cancellation fee (see Table 7). On the contrary, with an FTT, HF orders' aggressiveness is higher than the one of LF traders, especially in the recovery phase.

³³For a different view see the work of Mannaro et al (2008).

³⁴Notice that we ran the above experiments for other values of γ^H and different levels of *ftt*. However, simulation results are consistent with the ones presented in Table 8.

Table 8: The effect of different transaction tax levels on price volatility and flash crash statistics when $\gamma^H = 20$. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

$\gamma^H = 20$	ftt	σ_P	Number of flash crashes	Avg. duration of flash crashes
	0%	0.016 (0.002)	4.636 (0.398)	7.139 (0.484)
	0.05%	0.010 (0.000)	3.279 (0.271)	7.782 (0.843)
	0.5%	0.009 (0.000)	2.697 (0.249)	8.144 (0.800)
	1%	0.009 (0.000)	3.094 (0.306)	8.753 (0.690)
	10%	0.004 (0.000)	1.429 (0.107)	11.286 (1.147)
	50%	0.002 (0.000)	- -	- -




Figure 17: Market volatility as a function of minimum resting times duration. Values on the y-axis are averages across 50 independent Monte-carlo runs and are expressed in relation to the baseline value in the unregulated market.




Figure 18: Average number (solid line, left scale) and average duration of flash crash (dashed line, right scale) as a function of minimum resting times duration. Values on the y-axis are averages across 50 independent Monte-carlo runs and are expressed in relation to the baseline value in the unregulated market.

The above outcomes are explained by the different mechanisms through which cancellation fees and transaction taxes transmit their effects in markets. As we discussed in Section 3.3.3, a cancellation fee encourages HF traders to keep their orders in the book. As a result, orders' cancellation is reduced as well as orders' aggressiveness. In contrast, an FTT boosts order cancellation by increasing the required expected profit threshold to keep an order in the book. Without a transaction tax, a HF trader has the incentive to maintain the order in the book if the expected profit is non-negative $\pi_{j,t}^E \geq 0$. In our model, with a transaction tax, an order is kept in the book if $\pi_{j,t}^E \geq ftt \cdot D_{j,t}$. Thus the higher is the transaction tax rate, the larger is the amount of HF orders removed from the book in each trading session. However, sufficiently large amounts of order cancellations (as *e.g.*, it is the case for draconian tax rates) have the paradoxical effect of almost removing HF traders from the market, thus reducing volatility and leading flash crashes to vanish.

Overall, the above results cast doubts on the effectiveness of financial transaction taxes on HFT, especially if its validity is compared to the one other market-based measures such as cancellation fees (or command-and-control ones like minimum resting times). Indeed, besides exhibiting the same trade-off between stability and resilience already highlighted for the other policy measures, financial transaction taxes can achieve significant reductions in volatility and have some incidence on financial crashes only if they implemented at sufficiently high rates.

4 Concluding Remarks

We developed an agent-based model of a limit-order book (LOB) market based on Jacob Leal et al (2016) to analyze the effectiveness of a set of regulatory policies on market volatility, and on the occurrence and the duration of flash crashes. In the model, low-frequency (LF) traders interact with high-frequency (HF) agents. The former can switch between fundamentalist and chartist strategies. HF traders instead employ low-latency *directional* strategies to exploit the order book information released by LF agents. In addition, LF trading rules are based on *chronological* time, whereas HF ones are framed in *event* time, *i.e.*, the activation of HF traders endogenously depends on past price fluctuations. Finally, HF traders can endogenously cancel their orders from the book based on expected profits. In this framework, we analyzed via Monte-Carlo simulations, the impact of policies like *i.*) trading halt facilities (both ex-post and ex-ante designs); *ii.*) minimum resting times; *iii.*) order cancellation fees; *iv.*) transaction taxes. These policies have been proposed and implemented both in Europe and in the US to mitigate the possible damaging effects of HFT and to prevent flash crashes.

Computer simulations reveal that policies that hamper order cancellation by high-frequency traders, like the implementation of minimum resting times or cancellation fees lead to significant improvements in terms of lower market volatility and incidence of flash crashes. Also the introduction of a financial transaction tax, by discouraging HFT, can improve market stability, although the effectiveness of such a measure is much lower compared to policies targeting order cancellation, and effects are relevant only for high values of the tax. These results are all consistent with the remarks in Haldane (2014), who conjectures that the above set of policies are effective because they tackle the “race to zero” of HFT at source by imposing a speed limit on trading. At the same time, all these policies are characterized by a trade-off between market stability (in terms of lower volatility and number of flash crashes) and market resilience (in terms of longer recoveries from a crash). This trade-off emerges because of the positive role played by HFT in quickly restoring good liquidity conditions after a crash. Regulatory policies introduce important distortions in such a process, thereby contributing to lengthen the duration of price-recoveries. The beneficial impact of HFT on price resilience also underlies the results concerning the study of the impact of circuit breakers and, in particular, explain why ex-post circuit breakers have no effect on volatility and have a negative impact on the duration of flash crashes. In contrast, we find that ex-ante circuit breakers are very effective, as they markedly reduce price volatility and completely remove flash crashes.

Overall, our results suggest that regulatory policies can have quite complex effects

on markets populated by low and high-frequency traders. From the viewpoint of policy design, our analysis highlights in particular the importance of understanding the different transmission mechanisms through which the effects of regulatory policies unfold. Moreover, it points out the need of taking into account the fundamental dual role played by high-frequency traders. On the one hand, high-frequency trading can be the source of extreme events like flash crashes by placing aggressive sell orders and removing liquidity from the market. On the other hand, it can play a leading role in the recovery from the crash, by quickly restoring liquidity.

Our analysis could be extended in several ways. First, we could enlarge the set of policies considered, by including measures such as make/take fees, restrictions on tick size, position limits. Second, empirical evidence suggests that, under some circumstances, trading halts, meant to stabilize markets, can cause a magnet effect *i.e.*, price movements acceleration towards the preannounced limits as market participants alter their strategies and trade in anticipation of a market halt (*e.g.*, Cho et al, 2003; Hsieh et al, 2009). In this context, a future step towards a better understanding of the effectiveness of regulatory means probably requires to account for the strategic reaction of traders to regulation policies. Third, our policy results provide only a limited account of the overall impact of regulatory policies on market stability. Accordingly, extensions encompassing the analysis of other stability-related indicators (*e.g.*, mispricing and tail indices) could be envisaged. Some of these extensions, *e.g.*, an analysis of the effects of regulatory policies on mispricing, could also involve a careful analysis of how mispricing is affected by agents' trading strategies (both HFT and LFT ones), along the lines developed, for instance, in Westerhoff (2003b). Lastly, so far, we only have considered one asset market in the model. However, regulatory authorities should also focus on the linkages across markets, recognizing that some coordination is needed to ensure the effectiveness of regulatory interventions (see CFTC and SEC, 2010; Furse et al, 2011), especially in high frequency markets, where HF traders can rapidly process and profit from the information stemming from different exchanges (*e.g.*, Wah and Wellman, 2013). The May 6, 2010 highlighted, for instance, the importance of the interconnectedness of equities and derivatives markets.

Acknowledgments

We are grateful to Andrea Roventini, Giorgio Fagiolo, Sylvain Barde, Jean-Luc Gaffard, Charles Noursair, Olivier Scaillet, Nobuyuki Hanaki, Simone Alfarano and Francesco Saraceno for stimulating discussions and fruitful comments. We also thank the participants of the WEHIA conference in Sophia Antipolis, May 2015, to the FEBS conference in Nantes, June 2015, to the MAGECO 2015 workshop in

Paris, November 2015, and to the PFMC conference in Paris, December 2015, and to the Workshop on “Complex Evolving System Approach in Economics”, in Sophia-Antipolis, January 2016 where earlier versions of this paper were presented. The authors also acknowledge the financial support of the EU H2020 Project ‘Distributed Global Financial Systems for Society’ (DOLFINS) grant n. 640772. All usual disclaimers apply.

Appendix: Robustness Analyses

In this appendix, we discuss the results of some robustness analyses carried out on the model presented in the paper. We begin with the analyses concerning the hypotheses about activation of HF traders in the model. Next, we present results for scenarios where HF traders adopt different price setting rules.

4.1 High-Frequency Traders’ Activation

In the model, HF traders’ activation is based on past price variation, according to equation (7). More specifically, each HF trader has a fixed price threshold Δx_j , drawn from a uniform distribution with support bounded between η_{min} and η_{max} . In all the experiments discussed in the paper, we selected $\eta_{min} = 0$ and $\eta_{max} = 0.20$ (cf. Table 1) in order to have sufficient variability in activation thresholds across HF traders and to allow HF traders to activate in presence of both large and small price changes. In order to test the robustness of our main results, we conduct additional simulations assuming that HF traders only activate due to small price changes *i.e.*, whenever:

$$\left| \frac{\bar{P}_{t-1} - \bar{P}_{t-2}}{\bar{P}_{t-2}} \right| \leq \Delta x_j$$

and where Δx_j is randomly selected for each agent from a uniform distribution with support in the open interval $(0, 0.05)$. Note that this implies that no HF trader will get activated when the relative variation in price is in absolute value equal or larger than 5%.³⁵ Table 9 shows the results of these experiments for the baseline scenario used in the paper and for the scenario with the unregulated market but with HF traders activating only for small price changes.

As the table shows, the statistics about volatility and flash crashes are not significantly affected by the introduction of HF traders that activate only for small changes.

Moreover, we also tested the results of the experiment about the minimum resting time policy (see Section 3.3.2) when HF traders are only activated with small price changes. Table 10 shows that the results of these additional exercises are consistent with the main results of our model. For instance, increasing minimum resting times mostly brings flash crash duration up to a threshold above which flash crashes eventually disappear, a property that is in line with our results about price resiliency in the paper.

³⁵In the model, flash crashes are identified as drops in the asset price of at least 5% followed by a sudden recovery of at most 30 minutes.

Table 9: Market volatility (σ_P) and flash crashes statistics in the baseline scenario and in the scenario where HF traders activate only for small price changes and $\eta \in (0, 0.05)$. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

	σ_P	Number of flash crashes	Avg. duration of flash crashes
Baseline	0.016 (0.002)	4.636 (0.398)	7.139 (0.484)
HF traders activating for small price changes only	0.018 (0.001)	6.681 (0.597)	8.237 (0.616)

Table 10: The effect of minimum resting times on price volatility and duration of flash crashes in the scenario where HF traders get activated only for small prices changes, and $\eta \in (0, 0.05)$. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

	σ_P	Number of flash crashes	Avg. duration of flash crashes
1	0.013 (0.001)	6.750 (0.597)	12.177 (0.659)
3	0.010 0.000	3.220 0.266	21.447 0.569
5	0.009 0.000	2.794 0.274	20.533 0.721
10	0.010 (0.000)	3.195 (0.276)	14.470 (0.920)
20	0.011 (0.001)	5.081 (0.411)	9.145 (0.642)
40	0.002 (0.000)	- -	- -

Finally, Table 11 shows the results of introducing ex-post circuit breakers in this new scenario with HF traders' activation based only on small price changes (all other parameters are set at the same values as in the paper).

Table 11: The effect of ex-post circuit breakers on price volatility and duration of flash crashes. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

	σ_P	Avg. duration of flash crashes
No circuit breaker	0.018 (0.001)	8.237 (0.616)
ex-post circuit breaker	0.021 (0.002)	13.452 (0.445)

As the table shows quite starkly, changing the sign of the inequality in equation 7 has not effect whatsoever on the result that the introduction of a ex-post circuit-breaker lengthens the duration of a flash crash (see Section 3.3.1). Indeed, in presence of an ex-post circuit breaker, the average duration of the flash crash is increased exactly by the duration of the trading halt ($np = 5$).³⁶

4.2 High-Frequency Traders' Pricing Strategies

As we explained in the paper (see Section 3.2), the difference in price setting rules between high- and low-frequency traders is one of the key ingredients in the generation of high volatility and flash crashes in our model. On the one hand, we find that HF traders do provide liquidity to the market most of the time. Indeed, order aggressiveness ratios are very low in normal times *i.e.*, the regime where the market spends most of the time during our simulations (see Table 3 and the discussion in Section 3.2). On the other hand, the pricing strategy of HF traders may also occasionally generate large bid ask spreads, one of the two ingredients of flash crashes in our model.

In this section, we turn to investigate whether high volatility and flash crashes emerge also when HF traders use alternative pricing strategies. We begin by exploring the scenario where HF buyers and sellers' limit prices are respectively formed as follows:

$$P_{j,t} = P_{j,t}^{bid}(1 + \kappa_j) \quad P_{j,t} = P_{j,t}^{ask}(1 - \kappa_j), \quad (11)$$

where κ_j is drawn from a uniform distribution with support ($\kappa_{min} = -0.01, \kappa_{max} = 0.01$). In this way, HF traders set their prices by taking the best price of their market side (best bid for buyers, best ask for sellers). Moreover, with the chosen setup for κ_j ,

³⁶We also experimented with other regulatory policies used in the paper and also for the case where the Δx_j were drawn from a distribution with support in the open interval $(0, 0.2)$, and the results were similar to the one reported in this appendix.

they can also place passive orders in the market. We kept all the other parameters' values and decision rules as in the baseline. The results of this experiment is reported in the second row of Table 12. For the sake of comparison, this table also reports the values of market statistics in the baseline scenario (first row of the table). As the table reveals quite starkly, changing the pricing strategy of HF traders has no effect whatsoever on market volatility and the number of flash crashes, which are unchanged with respect to the baseline values. The modification only yields a small increase in the duration of flash crashes.

Furthermore, we repeated the above experiment by also assuming that HF traders activate for small price changes only (see also the previous section). The results of this last experiment are reported in the third row of Table 12. Interestingly, combining the pricing strategy of equation 11 with the small price activation destabilizes the market rather than making it more stable. Indeed, both the volatility and the number of flash crashes significantly rise compared to the baseline.

Finally, we also repeated the experiment on the minimum resting policy (cf. Section 3.3.2 of the paper) in the alternative pricing scenario for HF traders and by assuming that they activate for small price changes only. The results of this last experiment are reported in Table 13. Again, the table shows that the main policy results highlighted in the paper are robust to this alternative scenario. Increasing minimum resting times has the effect of lowering volatility and the number of flash crashes. In addition, the duration of flash crashes increases.³⁷

Table 12: Market volatility (σ_P) and flash crashes statistics in the baseline scenario and in alternative HFT pricing and activation scenarios. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

	σ_P	Number of flash crashes	Avg. duration of flash crashes
Baseline	0.016 (0.002)	4.636 (0.398)	7.139 (0.484)
HFT also with passive orders	0.016 (0.000)	4.000 (0.348)	11.657 (0.665)
HFT also with passive orders and activating for small price changes only	0.035 (0.001)	31.580 (1.432)	9.123 (0.297)

In the last set of robustness exercises presented in this section, we assume that HF traders use the same pricing strategy as LF traders. More precisely, we perform

³⁷The only change is represented by the fact that flash crashes now emerge also for the longest length of the minimum resting times policy.

Table 13: The effect of minimum resting times on price volatility and duration of flash crashes in the scenario where HF traders can also submit passive orders and get activated only for small prices changes, and $\Delta x_j \in (0, 0.05)$. Values are averages across 50 independent Monte-Carlo runs. Monte-Carlo standard errors in parentheses. (σ_P): price returns volatility.

	σ_P	Number of flash crashes	Avg. duration of flash crashes
1	0.024 (0.000)	21.120 (1.186)	11.541 (0.278)
3	0.018 (0.000)	5.980 (0.353)	12.183 (0.480)
5	0.017 (0.000)	4.286 (0.290)	13.276 (0.552)
10	0.014 (0.000)	2.691 (0.215)	12.756 (0.795)
20	0.013 (0.000)	2.564 (0.217)	17.674 (1.099)
40	0.007 (0.000)	1.238 (0.076)	10.762 (1.068)

experiments where the limit-order price of each HF trader is determined by:

$$P_{j,t} = \bar{P}_{t-1}(1 + \delta)(1 + z_{j,t}), \quad (12)$$

where $\delta > 0$ and $z_{j,t}$ measures the number of ticks away from the last market price \bar{P}_{t-1} and it is drawn from a Gaussian distribution with zero mean and σ_z^H standard deviation. We repeat the same experiments for different values of the price ticks standard deviation σ_z^H . Notice that the latter Gaussian distribution is centered around zero, which implies that in these additional experiments, the mode and the mean of the distribution of HF orders' prices is always very close to the last closing price and it is always the same as for LF traders.

The results of these experiments are reported in Table 14 below. Interestingly, the table shows that high volatility and flash crashes do arise in our model also when HF traders employ the same pricing strategy as LF traders, as long as the price-tick standard deviation of HF traders (σ_z^H) is enough larger than the one of LF traders ($\sigma_z^L = 0.01$ in our model, see Table 1 in the paper).

Table 14: Market volatility (σ_P) and flash crashes statistics in the scenario where HF traders adopt the same pricing strategy of LF traders and for varying levels of the variance of distribution of HF traders' price ticks (σ_z^H).

σ_z^H	σ_P	Number of flash crashes	Avg. duration of flash crashes
0.01	0.005 (0.000)	- -	- -
0.02	0.010 (0.000)	1.000 (0.000)	14.000 (1.103)
0.03	0.013 (0.000)	2.182 (0.195)	11.593 (0.824)
0.04	0.013 (0.000)	3.286 (0.287)	13.566 (0.900)
0.05	0.013 (0.000)	3.732 (0.349)	12.134 (0.836)
0.1	0.023 (0.001)	7.400 (0.684)	10.425 (0.524)

References

- Abhyankar A, Copeland L, Wong W (1995) Moment condition failure in high frequency financial data: evidence from the s&p 500. *Applied Economics Letters* 2(8):288–290
- Ackert L (2012) The impact of circuit breakers on market outcomes. Tech. rep., Economic Impact Assessment EIA9, Foresight, Government Office for Science, London (UK).
- Aldridge I (2013) High-frequency trading: a practical guide to algorithmic strategies and trading systems. John Wiley & Sons
- Alfarano S, Lux T, Wagner F (2010) Excess volatility and herding in an artificial financial market: Analytical approach and estimation. Paper No. 24719, MPRA
- Aloud M, Tsang E, Olsen R, Dupuis A (2012) A directional-change event approach for studying financial time series. *Economics: The Open-Access, Open-Assessment E-Journal* 6(2012-36):1–17
- Aloud M, Tsang E, Olsen R (2013) Modeling the fx market traders' behavior: An agent-based approach. In: Alexandrova-Kabadjova B, Martinez-Jaramillo S, Garcia-

- Almanza A, Tsang E (eds) *Simulation in Computational Finance and Economics: Tools and Emerging Applications*, Hershey, PA: Business Science Reference
- AMF (2010) High frequency trading: the application of advanced trading technology in the european marketplace. Tech. rep., Authority for the Financial Markets (AFM), available at: <http://www.afm.nl/layouts/afm/default.aspx /media/files/rapport/2010/hft-report-engels.ashx>
- Ané T, Geman H (2000) Order flow, transaction clock and normality of asset returns. *Journal of Finance* 55:2259–2284
- Angel JJ, Harris LE, Spatt CS (2011) Equity trading in the 21st century. *The Quarterly Journal of Finance* 1(01):1–53
- Apergis N (2014) The role of circuit breakers in the oil futures market. *Journal of Economics and Finance* pp 1–16
- Avellaneda M, Stoikov S (2008) High-frequency trading in a limit order book. *Quantitative Finance* 8(3):217–224
- Baron M, Brogaard J, Kirilenko A (2014) Risk and return in high frequency trading. Working paper, SSRN
- Bartolozzi M (2010) A multi agent model for the limit order book dynamics. *European Physical Journal B: Condensed Matter Physics* 78(2):265–273
- Bershova N, Rakhlin D (2013) High-frequency trading and long-term investors: A view from the buy-side. *Journal of Investment Strategy* 2(2):25–69
- Biais B, Woolley P (2011) High frequency trading. Manuscript, Toulouse University, IDEI
- Biais B, Foucault T, Moinas S (2015) Equilibrium fast trading. *Journal of Financial Economics* 116(2):292–313
- Bildik R, Gülay G (2006) Are price limits effective? evidence from the istanbul stock exchange. *Journal of Financial Research* 29(3):383–403
- Bollerslev T, Wright JH (2000) Semiparametric estimation of long-memory volatility dependencies: The role of high-frequency data. *Journal of econometrics* 98(1):81–106
- Breckenfelder JH (2013) Competition between high-frequency traders, and market quality. Available at SSRN 2264858
- Brewer P, Cvitanic J, Plott CR (2013) Market microstructure design and flash crashes: A simulation approach. *Journal of Applied Economics* 16(2):223–250

- Brock W, Hommes C (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22(8-9):1235–1274
- Brogaard J (2010) High frequency trading and its impact on market quality. Northwestern University Kellogg School of Management Working Paper
- Brugler J, Linton OB (2014) Single stock circuit breakers on the london stock exchange: do they improve subsequent market quality? Available at SSRN 2379029
- Budish EB, Cramton P, Shim JJ (2013) The high-frequency trading arms race: Frequent batch auctions as a market design response. Fama-miller working paper
- CFTC, SEC (2010) Findings regarding the market events of may 6, 2010. Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues p 104
- Chakraborti A, Toke IM, Patriarca M, Abergel F (2011) Econophysics review: I. empirical facts. *Quantitative Finance* 11(7):991–1012
- Chiarella C, He X (2003) Heterogeneous beliefs, risk, and learning in a simple asset-pricing model with a market maker. *Macroeconomic Dynamics* 7(4):503–536
- Cho DD, Russell J, Tiao GC, Tsay R (2003) The magnet effect of price limits: evidence from high-frequency data on taiwan stock exchange. *Journal of Empirical Finance* 10(1):133–168
- Clark PK (1973) A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* 41:135–155
- Clauset A, Shalizi CR, Newman ME (2009) Power-law distributions in empirical data. *SIAM review* 51(4):661–703
- Colliard JE, Hoffmann P (2013) Sand in the chips: Evidence on taxing transactions in an electronic market. Tech. rep., mimeo
- Cont R (1997) Scaling and correlation in financial time series: a study of *s&p* futures prices. arXiv preprint cond-mat/9705075
- Cont R (2001) Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance* 1(2):223–236
- Cont R, Potters M, Bouchaud JP (1997) Scaling in stock market data: stable laws and beyond. In: *Scale invariance and beyond*, Springer, pp 75–85
- Copeland TE, Galai D (1983) Information effects on the bid-ask spread. *the Journal of Finance* 38(5):1457–1469
- De Long BJ, Shleifer A, Summers LH, Waldmann RJ (1990) Noise trade risk in financial markets. *Journal of Political Economy* 98:703–738

- Demary M (2010) Transaction taxes and traders with heterogeneous investment horizons in an agent-based financial market model. *Economics-The Open-Access, Open-Assessment E-Journal* 4(8):1–44
- Diacogiannis GP, Patsalis N, Tsangarakis* NV, Tsiritakis ED (2005) Price limits and overreaction in the athens stock exchange. *Applied Financial Economics* 15(1):53–61
- Dupont DY, Lee GS (2007) Effects of securities transaction taxes on depth and bid-ask spread. *Economic Theory* 31(2):393–400
- Easley D, López de Prado M, O’Hara M (2012) The volume clock: Insights into the high frequency paradigm. *Journal of Portfolio Management* 39(1):19–29
- Ehrenstein G (2002) Cont–bouchaud percolation model including tobin tax. *International Journal of Modern Physics C* 13(10):1323–1331
- Fama EF (1989) Perspectives on October 1987, or, what did we learn from the crash? Center for Research in Security Prices, Graduate School of Business, University of Chicago
- Farmer JD (2002) Market force, ecology and evolution. *Industrial and Corporate Change* 11(5):895–953
- Farmer JD, Skouras S (2013) An ecological perspective on the future of computer trading. *Quantitative Finance* 13(3):325–346
- Farmer JD, Patelli P, Zovko II (2005) The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences of the United States of America* 102(6):2254–2259
- Foucault T, Röell A, Sandås P (2003) Market making with costly monitoring: An analysis of the soes controversy. *Review of Financial Studies* 16(2):345–384
- Franke R, Westerhoff F (2009) Validation of a structural stochastic volatility model of asset pricing. Christian-Albrechts-Universität zu Kiel Department of Economics
- Franke R, Westerhoff F (2011) Estimation of a structural stochastic volatility model of asset pricing. *Computational Economics* 38(1):53–83
- Franke R, Westerhoff F (2012) Structural stochastic volatility in asset pricing dynamics: estimation and model contest. *Journal of Economic Dynamics and Control* 36:1193–1211
- Fricke D, Lux T (2015) The effects of a financial transaction tax in an artificial financial market. *Journal of Economic Interaction and Coordination* 10(1):119–150
- Friederich S, Payne R (2015) Order-to-trade ratios and market liquidity. *Journal of Banking & Finance* 50:214–223

- Furse C, Haldane A, Goodhart C, Cliff D, Zigrand J, Houstoun K, Linton O, Bond P (2011) The future of computer trading in financial markets. Tech. rep., Working paper, Foresight, Government Office for Science, London (UK)., URL <http://www.bis.gov.uk/foresight/our-work/projects/current-projects/computer-trading>
- Giardina I, Bouchaud JP (2004) Volatility clustering in agent based market model. In: *The Complex Dynamics of Economic Interaction*, Springer, pp 171–196
- Goldstein MA, Kavajecz KA (2004) Trading strategies during circuit breakers and extreme market movements. *Journal of Financial Markets* 7(3):301–333
- Golub A, Keane J, Poon SH (2012) High frequency trading and mini flash crashes. Tech. Rep. 1212.6667, arXiv
- Gomber P, Haferkorn M (2013) High-frequency-trading. *Business & Information Systems Engineering* 5(2):97–99
- Gomber P, Haferkorn M, Zimmermann K (2015) Securities transaction tax and market quality—the case of france. *European Financial Management*
- Gopikrishnan P, Plerou V, Amaral LAN, Meyer M, Stanley HE (1999) Scaling of the distribution of fluctuations of financial market indices. *Physical Review E* 60(5):5305
- Gopikrishnan P, Plerou V, Liu Y, Amaral LN, Gabaix X, Stanley HE (2000) Scaling and correlation in financial time series. *Physica A: Statistical Mechanics and its Applications* 287(3):362–373
- Gorski A, Drozd S, Speth J (2002) Financial multifractality and its subtleties: an example of dax. *Physica A: Statistical Mechanics and its Applications* 316(1):496–510
- Greenwald BC, Stein JC (1991) Transactional risk, market crashes, and the role of circuit breakers. *Journal of Business* pp 443–462
- Griffith-Jones S, Persaud A (2012) Financial transaction taxes. Study prepared for—and presented to—the Committee on Economic and Monetary Affairs of the European Parliament
- Haberer M (2004) Might a securities transactions tax mitigate excess volatility? some evidence from the literature. *Some Evidence from the Literature* (April 2004) University of Konstanz Discussion Paper (04/06)
- Habermeier KF, Kirilenko AA (2001) Securities transaction taxes and financial markets. *IMF Working Papers* pp 1–29
- Haldane A (2014) The race to zero. Speech, Bank of England
- Hanson TA (2011) The effects of high frequency traders in a simulated market. Available at SSRN 1918570

- Hasbrouck J, Saar G (2013) Low-latency trading. *Journal of Financial Markets* 16(4):646–679
- Hau H (2006) The role of transaction costs for financial volatility: Evidence from the paris bourse. *Journal of the European Economic Association* 4(4):862–890
- Hayes R, Paddrik M, Todd A, Yang S, Beling P, Scherer W (2012) Agent based model of the e-mini future: application for policy making. In: *Proceedings of the 2012 Winter Simulation Conference, Winter Simulation Conference*, p 111
- Hendershott T, Jones CM, Menkveld AJ (2011) Does algorithmic trading improve liquidity? *Journal of Finance* 66(1):1–33
- Hommes C, Huang H, Wang D (2005) A robust rational route to randomness in a simple asset pricing model. *Journal of Economic Dynamics and Control* 29(6):1043–1072
- Hsieh PH, Kim YH, Yang JJ (2009) The magnet effect of price limits: A logit approach. *Journal of Empirical Finance* 16(5):830–837
- Jacob Leal S, Napoletano M, Roventini A, Fagiolo G (2016) Rock around the clock: An agent-based model of low-and high-frequency trading. *Journal of Evolutionary Economics* 26(1):49–76, DOI 10.1007/s00191-015-0418-4
- Johnson N, Zhao G, Hunsader E, Meng J, Ravindar A, Carran S, Tivnan B (2012) Financial black swans driven by ultrafast machine ecology. Tech. Rep. 1202.1448, arXiv
- Jones CM, Seguin PJ (1997) Transaction costs and price volatility: evidence from commission deregulation. *The American Economic Review* pp 728–737
- Kim KA, Rhee S (1997) Price limit performance: evidence from the tokyo stock exchange. *The Journal of Finance* 52(2):885–901
- Kirilenko A, Kyle A, Samadi M, Tuzun T (2011) The flash crash: The impact of high frequency trading on an electronic market. Available at SSRN 1686004
- Kirilenko AA, Lo AW (2013) Moore’s law versus murphy’s law: Algorithmic trading and its discontents. *The Journal of Economic Perspectives* pp 51–72
- Kirman A, Teyssiere G (2002) Microeconomic models for long memory in the volatility of financial time series. *Studies in Nonlinear Dynamics & Econometrics* 5(4):281–302
- Kodres LE, O’Brien DP (1994) The existence of pareto-superior price limits. *The American Economic Review* pp 919–932
- Lauterbach B, Ben-Zion U (1993) Stock market crashes and the performance of circuit breakers: Empirical evidence. *The Journal of Finance* 48(5):1909–1925

- Lavicka H, Lichard T, Novotny J (2014) Sand in the wheels or wheels in the sand? tobin taxes and market crashes. Working paper series cerge-ei, The Center for Economic Research and Graduate Education-Economic Institute, Prague
- Lee WB, Cheng SF, Koh A (2011) Would price limits have made any difference to the 'flash crash' on may 6, 2010. *The Review of Futures Markets* 19(special ifm issue):55–93
- Lin TC (2012) The new investor. *UCLA L Rev* 60:678
- Lux T, Marchesi M (2000) Volatility clustering in financial markets: A microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance* 3(4):675–702
- Mandelbrot B, Taylor M (1967) On the distribution of stock price differences. *Operations Research* 15:1057–162
- Mannaro K, Marchesi M, Setzu A (2008) Using an artificial financial market for assessing the impact of tobin-like transaction taxes. *Journal of Economic Behavior & Organization* 67(2):445–462
- Mantegna RN, Stanley HE, et al (1995) Scaling behaviour in the dynamics of an economic index. *Nature* 376(6535):46–49
- Maslov S (2000) Simple model of a limit order-driven market. *Physica A: Statistical Mechanics and its Applications* 278(3):571–578
- McCulloch N, Pacillo G (2011) The tobin tax: a review of the evidence. *IDS Research Reports* 2011(68):1–77
- Menkveld AJ (2013) High frequency trading and the new-market makers. *Quarterly Journal of Economics* 128(1):249–85
- Mike S, Farmer JD (2008) An empirical behavioral model of liquidity and volatility. *Journal of Economic Dynamics and Control* 32(1):200–234
- Mills TC (2001) Statistical analysis of high frequency data from the athens stock exchange. *Physica A: Statistical Mechanics and its Applications* 293(3):566–572
- Paddrik ME, Hayes RL, Todd A, Yang SY, Scherer W, Beling P (2011) An agent based model of the e-mini s&p 500 and the flash crash. Available at SSRN 1932152
- Pellizzari P, Westerhoff F (2009) Some effects of transaction taxes under different microstructures. *Journal of Economic Behavior & Organization* 72(3):850–863
- Pollin R, Baker D, Schaberg M (2003) Securities transaction taxes for us financial markets. *Eastern Economic Journal* pp 527–558

- Prewitt M (2012) High-frequency trading: Should regulators do more. Michigan Telecommunications and Technology Law Review 19:131
- Ranaldo A (2004) Order aggressiveness in limit order book markets. Journal of Financial Markets 7(1):53–74
- Roll R (1989) Price volatility, international market links, and their implications for regulatory policies. Journal of Financial Services Research 3(2-3):211–246
- Santoni GJ, Liu T (1993) Circuit breakers and stock market volatility. Journal of Futures Markets 13(3):261–277
- for Science GO (2012) What is the economic impact of the mifid rules aimed at regulating high-frequency trading? Tech. rep., Economic Impact Assessment, Foresight, Government Office for Science, London (UK)., URL https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/289062/12-1080-eia21-economic-impact-mifid-rules-high-frequency-trading.pdf
- SEC (2010) Concept release on equity market structure. Release No. 34-61358, january 14, 2010. Available at: <http://www.sec.gov/rules/concept/2010/34-61358.pdf>
- SEC (2011a) Sec announces filing of limit up-limit down proposan to address extraordinary market volatility. Press Release No. 31647-01, april 5, 2011. Available at: <http://www.sec.gov/news/press/2011/2011-84.htm>.
- SEC (2011b) Sec to publish for public comment updated market-wide circuit breaker proposals to address extraordinary market volatility. Press Release, september 27, 2011. Available at: <http://www.sec.gov/news/press/2011/2011-190.htm>.
- SEC (2012) Sec approves proposals to address extraordinary volatility in individual stocks and broader stock market. Tech. rep., Press Release SEC, URL <http://www.sec.gov/News/PressRelease/Detail/PressRelease/1365171482422>
- SEC (2014) Equity market structure literature review. part ii: High frequency trading. White Paper, march 18, 2014
- Selçuk F, Gençay R (2006) Intraday dynamics of stock market returns and volatility. Physica A: Statistical Mechanics and its Applications 367:375–387
- staff I (2010) Fair and substantial contribution by the financial sector. final report for the g-20
- Stamatiou TG (2008) Price limits, volatility, liquidity and abnormal returns: An event study from the athens stock exchange. Cambridge Scholars Publishing, Newcastle UT
- Subrahmanyam A (2012) Stock market circuit breakers. Tech. rep., Economic Impact Assessment EIA4, Foresight, Government Office for Science, London (UK).

- Subrahmanyam A (2013) Algorithmic trading, the flash crash, and coordinated circuit breakers. *Borsa Istanbul Review* 13(3):4–9
- Tobin J (1978) A proposal for international monetary reform. *Eastern economic journal* 4(3/4):153–159
- Umlauf SR (1993) Transaction taxes and the behavior of the swedish stock market. *Journal of Financial Economics* 33(2):227–240
- Vuorenmaa TA, Wang L (2014) An agent-based model of the flash crash of may 6, 2010, with policy implications. *Ssrn*
- Wah E, Wellman MP (2013) Latency arbitrage, market fragmentation, and efficiency: a two-market model. In: *Proceedings of the fourteenth ACM conference on Electronic commerce*, ACM, pp 855–872
- Westerhoff F (2003a) Heterogeneous traders and the tobin tax. *Journal of Evolutionary Economics* 13(1):53–70
- Westerhoff F (2003b) Speculative markets and the effectiveness of price limits. *Journal of Economic Dynamics and Control* 28(3):493–508
- Westerhoff FH (2004) Speculative dynamics, feedback traders and transaction taxes: a note. *Review of Economics* pp 190–195
- Westerhoff FH (2006) Technical analysis based on price-volume signals and the power of trading breaks. *International Journal of Theoretical and Applied Finance* 9(02):227–244
- Westerhoff FH (2008) The use of agent-based financial market models to test the effectiveness of regulatory policies. *Journal of Economics and Statistics (Jahrbuecher fuer Nationaloekonomie und Statistik)* 228(2+ 3):195–227
- Westerhoff FH, Dieci R (2006) The effectiveness of keynes–tobin transaction taxes when heterogeneous agents can trade in different markets: a behavioral finance approach. *Journal of Economic Dynamics and Control* 30(2):293–322
- Windrum P, Fagiolo G, Moneta A (2007) Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation* 10(2):8
- Yeh CH, Yang CY (2010) Examining the effectiveness of price limits in an artificial stock market. *Journal of Economic Dynamics and Control* 34(10):2089–2108
- Yeh CH, Yang CY (2013) Do price limits hurt the market? *Journal of Economic Interaction and Coordination* 8(1):125–153
- Zovko I, Farmer JD (2002) The power of patience: a behavioural regularity in limit-order placement. *Quantitative Finance* 2(5):387–392