Beyond Circuit Breakers: Leveraging Machine Learning for Market Stability

Grégoire KEMLER[*] and Nohade NASRALLAH[†] and Iryna VERYZHENKO[‡]

Abstract

To safeguard financial markets from crashes and bubbles, circuit breakers have been implemented as protective mechanisms. However, these tools have faced significant criticism. Recent research suggests that circuit breakers, rather than mitigating trading panic during extreme market events, may amplify volatility through a "magnet effect." Additionally, they are often criticized for relying solely on price variations, despite growing evidence that price movements can be anticipated through order book activity. In this study, we employ machine learning algorithms - specifically recurrent neural networks (RNNs) and random forests - to predict extreme market events that could trigger circuit breakers based on detailed order book and trading activity data. These advanced techniques excel at uncovering intricate patterns in large, high-frequency datasets that traditional methods may overlook. Our analysis utilizes microsecond-level data from Euronext Paris, capturing the full flux of messages and transactions from January 4, 2016, to December 28, 2016. The dataset is uniquely enriched with trader classifications provided by the French Market Authority, distinguishing among pure high-frequency traders (HFTs), mixed HFTs, and non-HFTs. Moreover, it identifies whether trades were executed or orders placed on a market member's own account, on behalf of clients, or under liquidity provision contracts. Our findings reveal that non-HFTs are the primary contributors to price instability leading to circuit breaker activations, particularly during downward trends. Conversely, HFTs, often criticized for destabilizing markets, exhibit stabilizing behaviors during circuit breakers, especially in post-halt scenarios. Our results advance the understanding of circuit breakers in modern financial markets and highlight the potential of machine learning to refine their design and effectiveness.

JEL Classification: G10, G14, C63

Keywords: Circuit Breakers, High-Frequency Trading, Artificial Intelligence, Machine Learning, Recurrent Neural Network, Random Forest

[*]École Nationale Supérieure d'Arts et Métiers, Paris, France

[†]CERIIM, Excelia Business School, Paris, France.

[‡]**Corresponding author, iryna.veryzhenko@lecnam.net, Conservatoire National des Arts et Métiers, Paris, France**

1 Introduction

Financial markets are dynamic systems characterized by their susceptibility to extreme events, such as market crashes and rapid price surges, which can destabilize the global economy (Lee and Schu, 2022). Circuit breakers have been implemented to provide a regulatory safeguard against such disruptions, offering temporary pauses in trading to stabilize volatile markets (Wang *et al.*, 2022). Their importance became evident following significant market disruptions, such as the 1987 crash, the 2010 Flash Crash (Subrahmanyam, 2013), flash crashes in Cryptocurrency exchanges in 2017, the Brexit referendum, and large intraday plunges of the DJIA in February 2018 (Sifat and Mohamad, 2020).

Yet, the implementation of circuit breakers remains a contentious topic as persistent criticism pinpointed some unintended consequences such as delayed price discovery and exacerbation of volatility through the magnet effect. The rise of algorithmic trading, particularly high-frequency trading (HFT), has further complicated market behaviors, demanding innovative approaches to managing extreme events (Wang *et al.*, 2022). In this regard, the integration of machine learning and big data analytics present a promising avenue for addressing these challenges (Yuan, 2024). Particularly, deep neural networks (DNNs) are revolutionizing financial market analysis with innovative methods to predict time series data and performance forecasting (Kolte *et al.*, 2023).

Existing studies offer mixed findings on the effectiveness of circuit breakers. While some research highlights their role in reducing panic and stabilizing markets others argue that circuit breakers may amplify volatility through the magnet effect (Wang *et al.*, 2022). The issue has transcended geographical limits with the increasing prevalence of cross-border market integrations, particularly within Europe (Brennan, 1986; Chowdhry and Nanda, 1998). As a result, there have been calls to harmonize regulatory measures across exchanges to prevent the migration of harmful order flows and maintain market liquidity (Chen *et al.*, 2024). Furthermore, the empirical focus has often overlooked the distinct behaviors of different trader categories – HFTs, mixed HFTs, and non-HFTs – during circuit breaker activations. Additionally, although machine learning has shown promise in financial forecasting, its potential to enhance circuit breaker calibration and prediction remains underexplored. This study bridges these gaps by combining granular order book data with machine learning techniques to analyze trader behaviors and predict circuit breaker triggers.

The Euronext Paris market, one of Europe's largest and most technologically advanced trading platforms, provides a fertile ground for examining circuit breakers. In 2016, significant geopolitical and financial events, such as the Brexit referendum and the collapse of Italy's Monte dei Paschi di Siena bank, resulted in unprecedented volatility. These events triggered multiple circuit breaker activations, highlighting the critical role of these mechanisms in stabilizing markets. The choice of Euronext Paris and 2016 as the focal point of this study is justified by the rich dataset of extreme market events, enabling a robust evaluation of circuit breaker efficacy in a high-stakes environment.

This study employs microsecond-level data from Euronext Paris, focusing on 125 automated circuit breakers triggered on french equities between January 4, 2016, and December 28, 2016. The data, sourced from Eurofidai's BEDOFIH database, provides detailed insights into the activities of pure HFTs, mixed HFTs, and non-HFTs. While previous studies have relied on proxies to estimate high-frequency trading activity (Hendershott *et al.*, 2011; Riordan and Storkenmaier, 2012), our dataset contains the full flux of messages and transactions, labeled by the French Market Authority to indicate the profiles of traders. Furthermore, our data allows us to distinguish whether a trade was executed or an order placed on a market member's own account, on behalf of clients, or in relation to a liquidity provision contract, among other scenarios. This enables a uniquely detailed analysis of the trading dynamics across different trader categories. Additionally, by employing recurrent neural networks (RNNs) and random forest algorithms, this study identifies the trader categories most influential in triggering circuit breakers and evaluates their contributions to market instability. The methodological approach integrates machine learning techniques to predict extreme events and refine circuit breaker mechanisms, addressing both theoretical and practical dimensions.

This paper makes several critical contributions to the literature. First, it provides a granular analysis of trader behaviors before, during, and after circuit breaker activations, distinguishing between HFTs, mixed HFTs, and non-HFTs. Second, it leverages machine learning to enhance the predictability of extreme market events, demonstrating the utility of advanced algorithms in regulatory applications. Third, by focusing on the high-stakes environment of Euronext Paris in 2016, the study offers unique insights into the effectiveness of circuit breakers during periods of geopolitical and financial instability. These contributions extend the discourse on market regulation, technological advancements, and the interplay between the two.

The results reveal that non-HFTs are the most significant contributors to price instability leading to circuit breaker activations, particularly during downward trends. Surprisingly, HFTs, often criticized for destabilizing markets, demonstrate stabilizing behaviors during circuit breakers, particularly in post-halt scenarios. Machine learning models achieve high accuracy in predicting circuit breaker triggers, with RNNs achieving a recall rate of 82.59%, indicating their effectiveness in identifying extreme events. However, both RNNs and random forests exhibit moderate precision, reflecting the challenge of minimizing false positives in volatile markets.

The remainder of this paper is organized as follows: Section 2 reviews the existing literature on circuit breakers, highlighting their theoretical foundations, empirical evidence, and gaps in research. Section 3 outlines the data sources and methodology, detailing the machine learning models and their implementation. Section 4 presents the results, focusing on trader behaviors and the predictive performance of machine learning algorithms. Section 5 discusses the findings in the context of regulatory implications and market stability. Finally, Section 6 concludes with policy recommendations and avenues for future research.

2 Literature Review

Circuit breakers are regulatory tools used to stabilize financial markets during periods of extreme volatility. In the European Union (EU), their implementation is governed under the Markets in Financial Instruments Directive II (MiFID II) framework. Article 48 mandates the use of mechanisms like trading halts and price limits to prevent extreme price fluctuations and ensure orderly markets (Lee and Schu, 2022). The Markets in Financial Instruments Directive II (MiFID II) framework in the European Union mandates circuit breakers, requiring trading venues to incorporate mechanisms calibrated for asset-specific liquidity and volatility profiles (ESMA, 2015).

The Paris Bourse operates under this regulatory framework, utilizing circuit breakers to manage volatility during significant market events like the 2016 Brexit referendum. On this occasion, trading halts on Euronext helped mitigate panic selling and facilitated the repricing of assets. As such, circuit breakers acted as a buffer, providing market participants with the time to reassess and mitigate abrupt sell-offs, thus stabilizing the CAC 40 index (Sifat and Mohamad, 2020). However, the global landscape shows variance in circuit breaker calibration, underscoring the need for dynamic and data-driven frameworks.

2.1 Theoretical and Empirical Foundations of Circuit Breakers

The theoretical basis for circuit breakers stems from the belief that allowing a pause in trading can alleviate information asymmetry, reduce order imbalances, and enhance liquidity (Greenwald and Stein, 1991). These mechanisms temporarily halt trading when predefined thresholds are breached. They are designed to achieve three primary objectives: (1) provide a cooling-off period during turbulent market conditions, (2) prevent panic selling, and (3) maintain orderly markets by allowing price discovery (GROSSMAN and MILLER, 1988).

Proponents of circuit breakers emphasize their ability to reduce panic and stabilize markets during periods of stress. For instance, Brennan (1986) highlights their role in lowering margin requirements and transaction costs, while Chowdhry and Nanda (1998) emphasize their function in limiting daily losses and preventing excessive speculation. Li *et al.* (2021) find that circuit breakers promote risk-sharing among investors, enabling a more equitable distribution of market risks.

However, this theoretical consensus is not universal. Critics point to the magnet effect, where circuit breakers inadvertently accelerate price movements toward the trigger points as traders rush to execute orders before a halt. Subrahmanyam (1994) demonstrates that this effect can increase price variability, reduce price efficiency, and destabilize markets. Goldstein and Kavajecz (2004) report liquidity reductions on the second day after trading halts. Chen *et al.* (2024) corroborate this, highlighting how circuit breakers in falling markets exacerbate volatility by encouraging concentrated trading activity near the thresholds. Jian *et al.* (2020) find that as thresholds approach, the likelihood of price jumps and trading surges increases, supporting the magnet effect hypothesis in the Chinese market. A circuit breaker often reduces the overall stock price level and significantly impacts its behavior. Specifically, as the price nears the circuit breaker threshold, volatility increases sharply, which in turn heightens the likelihood of triggering the mechanism–an effect commonly referred to as the magnet effect (Chen *et al.*, 2024). Subrahmanyam (2013) demonstrates that strategic trading near circuit breaker thresholds can precipitate unnecessary halts, amplifying market instability. Empirical evidence from the Tokyo Stock Exchange (Kim and Rhee, 1997) and the Chinese stock market (Wang *et al.*, 2019) supports this view, showing that trading activity intensifies as prices approach the limits, undermining the stabilizing intent of circuit breakers. The circuit breaker failed to act as a time-out buffer and instead triggered a rapid onset of the magnet effect, increasing the likelihood of stock prices reaching their limits during the market crash (Li *et al.*, 2021; Wong *et al.*, 2020).

2.2 The Role of Machine Learning in Predicting Circuit Breaker Activation

Machine learning has emerged as a powerful tool for predicting market crashes and optimizing circuit breaker mechanisms. Unlike traditional models, machine learning algorithms can process vast amounts of historical and real-time data, identifying complex, nonlinear patterns associated with market stress.

Furthermore, the integration of ML into financial market regulation has provided new avenues for enhancing the predictive and adaptive capabilities of circuit breakers. ML models, particularly deep neural networks (DNNs), can process high-frequency data to identify patterns and forecast market stress with exceptional accuracy (Boonpan and Sarakorn, 2025). In the context of our study, incorporating machine learning into the calibration of circuit breakers could enhance their efficacy by addressing challenges such as the magnet effect. For example, models trained on historical data could identify the behavioral patterns of market participants that precede trading halts. This predictive capability allows regulators and exchanges to implement proactive measures, such as adjusting threshold levels or introducing dynamic halts based on real-time conditions.

Many studies showcase the beneficial effects of ML. By analyzing liquidity, volatility, and order flow data, ML models can predict the likelihood of circuit breaker activation with high accuracy.

For instance, Boonpan and Sarakorn (2025) utilized advanced data preparation techniques, such as trend scanning and piecewise aggregate approximation, to enhance DNN-based predictions of stock price directions. Gao *et al.* (2024) employed agent-based models augmented with ML to simulate flash crash scenarios, offering insights into conditions that precipitate such events. Jian *et al.* (2020) demonstrated how ML models could mitigate the magnet effect by preemptively identifying high-risk scenarios. Additionally, agent-based simulations combined with ML algorithms provide a controlled environment for testing regulatory responses to hypothetical market shocks (Leal and Napoletano, 2019).

By leveraging real-time data, ML can dynamically adjust circuit breaker thresholds, addressing criticisms related to static calibrations. These advancements align with regulatory calls for adaptive and evidence-based calibration frameworks, as outlined in MiFID II and ESMA guidelines.

3 Data Description

Our data is provided by Eurofidai Base Européenne de Données Financières à Haute Fréquence (BEDOFIH). It covers large variety of stocks listed on Euronext from January 4, 2016 to December 28, 2016. Since January 2012, Eurofidai launched the Base Européenne de Données Financières à Haute Fréquence (BEDOFIH) project that aims to create a European intra-day financial database. One notable characteristic of this data is its granularity, recorded at the microsecond level. Each message within the database is assigned a category based on its owner. The Autorité des Marchés Financiers (AMF) has identified three categories of market participants: pure-HFTs, investment banks engaged in HFT activities (MIX), and all other remaining traders (Non-HFTs)[1].

Moreover, each trader is required to flag every order based on specific criteria, such as:

1. Own account or own account for client facilitation (OWN).
2. Own account of an affiliate or when operating from a parent company of the stock (PARENT).
3. Account of a third party or client account (CLIENT).
4. Orders submitted pursuant to a liquidity provision agreement (MM).
5. Orders submitted for retail liquidity provider (RLP) or retail matching facility (RMO).

Our analysis centers on the French equity market. Trading in the French stock market takes place on the Euronext Paris, which has been operating as a Regulated Market since 2009. The trading environment is computerized and utilizes a limit order system on the UTP (Universal Trading

[1] According to the AMF, there are 10 to 20 pure-HFT, 10 to 20 IB-HFT and 100 to 150 non-HFT operating on Euronext Paris.

Platform). Due to the increasing automation of stock markets, the time interval for submitting, modifying, or cancelling orders on the marketplace has become significantly shorter.

Euronext adheres to a general philosophy of not halting trading except under extreme circumstances. To maintain orderly markets, they employ a comprehensive set of safeguards designed to prevent disruptions and detect unusual events. These protection mechanisms are regularly reviewed and approved by regulators to ensure ongoing improvements in safety measures across all their regulated markets. To avoid disruptions, Euronext's market protection mechanisms include rejecting aberrant orders (e.g., unusually large in size or price) before they enter the market. They also have the ability to halt trading on specific instruments affected by single orders that could cause market disruption, as well as to prevent significant price fluctuations using dynamic and static collars.

Real-time algorithm-powered alerts are in place to monitor the markets and ensure fair and orderly trading, with investigation and resolution as necessary. Dynamic and static collars, based on the reference price that changes throughout the trading day after each trade, serve as the first and second layers of Euronext's circuit-breaker mechanism. Reservation periods during collars help mitigate the impact of unexpected price movements, allowing sufficient time for market participants to review orders and investment decisions before trading resumes.

On Euronext, a circuit breaker is recorded as a specific event with a series of states and their exact timing. The 'Halted instrument' state occurs when the static or dynamic price range is reached or breached, causing a temporary trading interruption. Unlike traditional trading halts, a circuit breaker does not fully suspend all trading activities but instead triggers an unscheduled call auction stage called 'Delayed opening.' During the call phase, orders can be submitted, modified, or canceled, and indicative prices and volumes are displayed. The purpose of this temporary trading suspension is to give market participants a few minutes to evaluate new information, reconsider their interests, and remove any erroneous orders. The minimum duration of a trading halt is 3 minutes. If the potential execution price remains outside the predetermined acceptable range, the auction can be extended. The 'Go to open' state occurs when the opening price is determined, and 'Start continuous trading' marks the moment when continuous trading resumes. This method of recording a circuit breaker allows for a reconstruction of all market events leading up to the circuit breaker and order book activity during the trading halt.

The descriptive statistics for our sample are presented in Table 1.

Figure 1 illustrates the companies listed on Euronext in 2016 whose stocks were halted due to circuit breaker events. Their price distribution is right-skewed, with the majority of stocks falling within the lower price range (0-30). A noticeable decline in frequency is observed as prices rise. The 0-20 range shows the highest frequency, suggesting that most halted stocks were relatively low-priced.

	HFT			Mix			Non-HFT		
	mean	median	sd	mean	median	sd	mean	median	sd
%NbOfTradesAtOpening	8.6302	7.3621	6.0885	47.1037	47.2393	9.9121	44.2662	45.3988	10.6580
%NbOfTradesAtContinuous	37.1946	37.9208	5.5548	53.9164	54.4951	4.9283	8.8890	7.5842	5.9310
%VolumeAtOpening	32.7510	33.5307	7.7573	78.8883	80.8941	10.2023	35.1900	31.3304	14.4435
%VolumeContinous	55.2593	56.7006	7.2372	85.2437	86.6005	5.1813	16.1757	13.7992	8.5683
TotalMarket	7.7838	7	4.2891	88.8378	84	27.9826	205.2973	175	133.1183
TotalLimit	76510.4324	71669	17467.9366	76862.7838	74133	18182.6766	2979.8108	2877	936.4472
TotalLimitMarket	0	0	0	0.4865	0	0.5588	50.1622	37	58.139
ExecutedMarket	2.7297	3	1.8507	40.8649	39	18.7026	68.5676	60	37.8649
ExecutedLimit	2324.8649	2236	798.325	3847.5135	3485	1826.4959	422.2162	313	268.6831
ExecutedMarketLimit	0	0	0	0.2432	0	0.435	13.6486	13	6.4602

Table 1: Summary statistics of different categories of traders. These statistics are measured based on the data of 102 equities companies listed on Euronext in the period from 04/01/2016 to 30/12/2016. %NbOfTradesAtOpenning = Nb of trades of a certain category of traders/nb of total trades realized during opening session. %NbOfTradesAtCotinous = Nb of trades of a certain category of traders/nb of total trades realized during continuous session.

Notable spikes at price levels of 60 and 100 indicate the presence of outliers. Lower-priced stocks may exhibit higher volatility, making them more susceptible to triggering circuit breakers. Among the largest companies affected by circuit breakers that year, BNP Paribas faced a significant legal challenge on April 29, 2016, when a class-action lawsuit was filed by a group of Sudanese refugees. The plaintiffs accused BNP Paribas of providing financial assistance to the Sudanese government, allegedly facilitating human rights violations, including genocide, in the Darfur region. On November 22, 2016, Vinci, a leading French construction and concessions company, saw a sharp decline in its stock price following the spread of a fraudulent press release. The false statement claimed that Vinci would restate its 2015 and 2016 financial reports and had dismissed its CFO. This misinformation rapidly circulated, causing Vinci's stock price to drop by more than 18% before trading was suspended. On June 24, 2016, Renault SA also experienced significant price volatility, triggering Euronext's circuit breaker. As a major European automaker, Renault was directly impacted by the market's reaction to the Brexit vote, with its stock price fluctuating sharply due to investor concerns about the economic and regulatory consequences of the UK's decision to leave the EU.

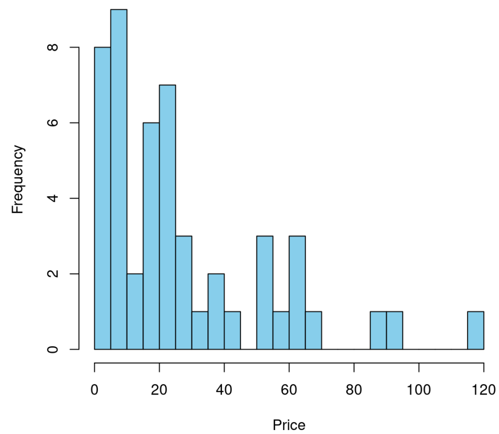

Figure 1: **Frequency of Opening Prices on the Day of Circuit Breaker Activation.**

The histogram 2 reveals two distinct periods of heightened circuit breaker activity on Euronext Paris in 2016, each corresponding to major geopolitical and financial events. A significant spike in circuit breakers occurred in late June, closely tied to the Brexit Referendum, which triggered extreme volatility and uncertainty across European markets. Earlier in the year, another cluster of circuit breakers was observed, primarily driven by the default of the Italian bank Monte dei Paschi di Siena. This event had a cascading effect on French banking and financial companies, exacerbating market instability. These findings highlight how systemic risks and major geopolitical events can lead to abrupt market dislocations, necessitating the activation of circuit breakers.

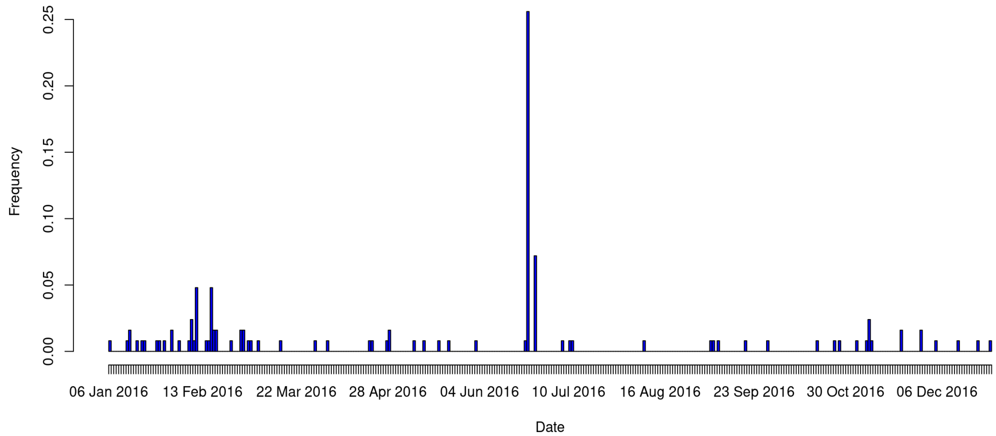

Figure 2: **Seasonal Patterns in Circuit Breaker Activations**

4 Traders' Activities Leading to Automated Circuit Breakers

To study the reaction of different categories of traders on stressful situation which leads to circuit breaker and their reaction on trading halts, we focus on a sample of 125 automated circuit breakers triggered on the equities listed on Euronext from January 4, 2016 to December 28, 2016. We report that 42 breaking episodes were provoked by significant daily price drop (called "negative circuit breaker"), and 83 of circuit breakers were provoked by significant daily price rise (called "positive circuit breaker"). We find that non-HFTs significantly increase their usage of market orders before-, during- and after trading pause. The ratio of market-to-limit registers about 10% for this category of traders. They do not have the same abilities as HFTs to monitor and manage their outstanding limit orders in highly volatile market, so to prevent adverse selection risk they heavily rely on market orders during extreme market events. More surprisingly, we find that fast traders also increase their trading aggressiveness and increase the usage of market orders during a circuit breaker. Market-to-limit ratio of pure-HFTs is three times higher during the halting and delayed opening compared to the normal market conditions. HFTs reduce their liquidity provision via limit orders during a circuit breaker.

To study the changes in liquidity provision and consumption of different categories of traders during relatively calm periods and particular turbulent periods, we study the monetary net trade imbalance in Table 2. That is the difference between the funds invested to buy transactions and funds gained as result of sell transactions. Negative net imbalance of a trading category during a crash indicates that it contributes to price drop; a positive net imbalance during a crash indicates that this category of

	Upward trend								
	pure-HFT			mix-HFT			non-HFT		
	mean	median	sd	mean	median	sd	mean	median	sd
5 minutes prior	36,561.07	0.00	101,959.4	79,861.22	4,407.82	217,561.3	121,280.20	10,671.61	356,494.9
halted intrument	-12,065.39	0.00	150,370.3	-85,797.16	0.00	529,709.1	-51,133.22	0.00	355,598.2
delayed opening	-44,583.87	-904.05	263,830.7	-140,474.19	-1,526.16	1,171,636.2	-388,644.11	-2,972.85	10,21,570.2
5 minutes after	10,460.15	4,621.21	382,539.9	-220,021.15	15,708.76	1,510,767.9	-247,376.83	-4,565.18	1,418,433.3
	Downward trend								
	pure-HFT			mix-HFT			non-HFT		
	mean	median	sd	mean	median	sd	mean	median	sd
5 minutes prior	-142,725.1	-8,608.923	386,429.6	-474,141.6	-41,055.007	15,85,882.3	-384,056.2	-94,598.257	755417.9
halted intrument	-1,219.075	0.00	55,549.34	-149,697.718	0.00	820,600.74	-40,225.256	0.00	303,384.10
delayed opening	-11,859.8	0.00	155,826.7	-148,345.3	-4,473.592	563,440.2	-168,807.0	-4,473.592	623,431.8
5 minutes after	53,510.91	-3,438.496	294,727.6	-266,183.86	-19,534.643	958,306.3	40,706.74	-24,104.517	870,640.0

Table 2: Net Positions are calculated as the sum of buy volume consumed and buy volume provided, minus the dollar sell volume consumed and sell volume provided. The statistics are derived from 83 circuit breakers triggered by upward trends and 42 circuit breakers triggered by strong downward trends on Euronext in 2016.

traders contribute to a market stabilization and price recovery. Analysing market participants' quoting behavior during trading pauses shows that non-HFTs increase liquidity consumption and create the strongest liquidity pressure into the direction of underlying trend 5 minutes prior- and during the first stage of a typical automated circuit breaker.

At the same time, they significantly reduce their liquidity pressure during the period of delayed opening and contribute to the price recovery when the continuous trading is renewed. Hence, non-HFTs start playing a role of liquidity suppliers and efficiently use trading pause to moderate price movements. The trading pause creates a layer of protection for market participants willing to trade against the price movement and initiate a trend reversal, as they are protected from adverse selection during the circuit breaker.

The behavior of HFTs differs with respect to the direction of the trend. They contribute to market correction during a typical positive circuit breakers (at delayed opening) triggered by a strong upward trend. Pure-HFTs similarly to other categories provide liquidity into the opposite direction of the positive underlying trend. However, only mix-HFTs continue creating selling-pressure in case of market crash during the last stages of circuit breaker. Mix-HFTs tend to be substantial sellers during the price fall and after the circuit breaker is triggered. Even 5 minutes after reopening, average net positions of mix-HFTs is negative reaching on average -200 000 euros. It means that mix-HFTs tend to initiate more trades into direction of an average crash rather than they stabilize prices. These findings suggest that a circuit breaker may be insufficient to dissuade all traders from trend following and prevent the crash amplification.

	Normal			Incident			
	Mean	Median	Sd	Mean	Median	Sd	p-value
Nb modif pure-HFT	3309.2189	506.2235	6294.2290	5884.2068	1716.2727	10510.9105	0.0224**
Cancel ratio pure-HFT	0.6719	0.8407	0.3332	0.5906	0.7695	0.3413	0.0643*
Buy/sell pure-HFT	2.0467	1.0330	4.9404	2.4480	1.0349	5.6994	0.5618
Market/limit pure-HFT	0.0347	0.0000	0.2099	0.0318	0.0000	0.1284	0.8973
NbTradesLinked pure-HFT	0.1841	0.0681	0.4878	0.2353	0.1176	0.4047	0.3802
Trade volume pure-HFT	2386.2568	1834.7617	2354.0556	2443.2073	2458.4237	2246.3063	0.8488
Nb trades pure-HFT	33.8034	7.0000	63.5892	93.3306	27.0000	201.8748	0.0024***
Nb modif mix-HFT	10477.2727	3498.2515	15251.2792	11352.8977	3380.2985	20324.8010	0.7067
Cancel ratio mix-HFT	0.5793	0.6421	0.2501	0.5947	0.6620	0.2342	0.6241
Buy/sell mix-HFT	1.5046	0.9556	2.0873	1.5368	0.9695	3.2499	0.9275
Market/limit mix-HFT	0.1369	0.0000	0.5201	0.2537	0.0006	0.8138	0.1870
NbTradesLinked mix-HFT	0.2040	0.1221	0.2613	0.4259	0.2269	0.6299	0.0005***
Trade volume mix-HFT	2415.6540	2421.1655	1863.8331	2981.3059	3006.5078	2569.3283	0.0526*
Nb trades mix-HFT	32.6667	14.0000	59.1762	114.6942	34.0000	287.6880	0.0026***
Nb modif non-HFT	9732.9432	680.9231	24669.6050	8001.5353	1498.5135	20497.9252	0.5572
Cancel ratio non-HFT	0.3237	0.2727	0.2627	0.3109	0.2909	0.1861	0.6658
Buy/sell non-HFT	6.2106	1.3333	28.1359	1.8454	1.0732	2.3156	0.0971*
Market/limit non-HFT	0.1604	0.0096	0.5212	0.2101	0.1343	0.2985	0.3694
NbTradesLinked non-HFT	0.7812	0.5789	0.7464	1.0387	0.9354	0.8455	0.0133**
Trade volume non-HFT	2515.5250	2164.6814	4066.9374	2998.0587	2496.0603	3617.9166	0.3351
Nb trades non-HFT	25.6410	5.0000	45.6764	131.4793	34.0000	257.2837	0.0000***

Table 3: This table compares order book dynamics and trading activity across various trader categories during the 10 minutes preceding a circuit breaker trigger. The statistics labeled *normal conditions* correspond to the same 10-minute interval on typical trading days without circuit breaker activation.

Nb modif represents the average number of modified orders during the period. *Cancel ratio* refers to the ratio of canceled limit orders to submitted limit orders. *Buy/sell* indicates the ratio of bid orders to ask orders, while *Market/limit* shows the ratio of market orders to limit orders during the same period. *Order-to-trades* measures the average number of trades associated with each order, reflecting the number of trades an order is involved in. *Trade volume* represents the average trade volume for each trader category, and *Nb trades* denotes the average number of trades executed by each trader category during the period.

Significance levels are denoted as follows: ***: $p-value < 0.01$ (highly significant), **: $0.01 \leq p-value < 0.05$ (moderately significant), *: $0.05 \leq p-value < 0.1$ (weakly significant).

Table 3 presents the changes in order book dynamics and trading activity across different trader categories in the 10 minutes preceding a circuit breaker trigger. A notable rise is observed in the number of modifications made by pure-HFTs during circuit breaker events compared to normal conditions. There is also a weakly significant decrease in transaction annulment rates by pure-HFT during incidents. All trading categories exhibit a significant increase in both the number of trades and trading volume during this period. Transactions involving mix-HFTs and non-HFTs surge sixfold.

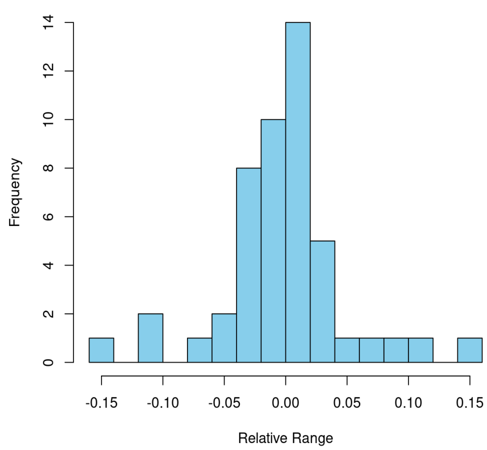

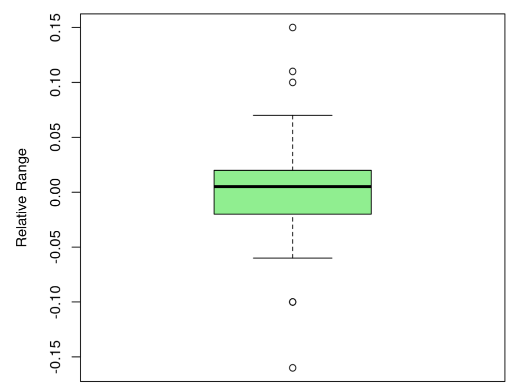

Figure 3: **Distribution of Relative Range 10 minutes prior to a circuit breaker**. To capture the underlying trend leading to a circuit breaker activation, we compute the relative range. First, we determine the indexes of the minimum price $P_i^{\min}$ and maximum price $P_j^{\max}$ within the interval $[P_{t-10}, P_t]$, where t is the moment of a circuit breaker activation, and $t-10 \leq i \leq t$ and $t-10 \leq j \leq t$. If $i > j$, then Relative Range $= \frac{P_i^{\min} - P_j^{\max}}{P_j^{\max}}$. If $i < j$, then Relative Range $= \frac{P_j^{\max} - P_i^{\min}}{P_i^{\min}}$. This histogram and boxplot summarize the distribution of the relative range. The mean is 0.0023, the median is 0.005, the standard deviation is 0.0505, the minimum value is -0.16, and the maximum value is 0.15.

The regression results in Table 4 examine the relationship between trade activity and the relative price range, a measure of price movement leading up to a circuit breaker activation. This analysis is presented for three trader categories: pure-HFT, mix-HFT, and non-HFT, with each category examined over two intervals: 10 minutes prior to the circuit breaker (CB) activation and 10 minutes after reopening.

For pure-HFT traders, the variable $NbTrades_{\text{Client}}$, which represents the number of trades executed on behalf of clients, shows a positive and statistically significant relationship with the relative

		10 minutes prior CB		10 minutes after reopening	
Category	Variable	Estimate	p-value	Estimate	p-value
pure-HFT	$NbTrades_{Client}$	1.066	0.0206*	0.619	0.2292
	$NbTrades_{Own}$	0.784	0.1128	-0.699	0.0691.
	$NbTrades_{Liquidity}$	0.706	0.6701	-0.442	0.4985
	$NbTrades_{RMO}$	-0.171	0.7502	-0.147	0.7154
	$NbTrades_{Others}$	0.603	0.6921	-0.730	0.1870
mix-HFT	$NbTrades_{Client}$	-0.064	0.8751	0.564	0.1523
	$NbTrades_{Own}$	-1.831	0.0122*	0.613	0.1126
	$NbTrades_{Liquidity}$	-0.170	0.9289	-0.370	0.5036
	$NbTrades_{RMO}$	-0.578	0.0819.	-0.001	0.9977
	$NbTrades_{Others}$	-0.169	0.9080	1.959	0.0602.
non-HFT	$NbTrades_{Client}$	-0.808	0.1228	-0.468	0.2713
	$NbTrades_{Own}$	0.860	0.0210*	-0.238	0.5002
	$NbTrades_{Liquidity}$	-1.102	0.0204*	0.335	0.2342
	$NbTrades_{RMO}$	0.648	0.0112*	0.231	0.6364
	$NbTrades_{Others}$	0.262	0.7008	-1.125	0.0791.
R^2		0.7415		0.4421	

Table 4: Regression Results of Relative Range on Trade Activity Across Trader Categories. To capture the underlying trend leading to a circuit breaker activation, we compute the relative range. First, we determine the indexes of the minimum price $P_i^{\min}$ and maximum price $P_j^{\max}$ within the interval $[P_{t-10}, P_t]$, where t is the moment of a circuit breaker activation, and $t - 10 \leq i \leq t$ and $t - 10 \leq j \leq t$. If $i > j$, then Relative Range $= \frac{P_i^{\min} - P_j^{\max}}{P_j^{\max}}$. If $i < j$, then Relative Range $= \frac{P_j^{\max} - P_i^{\min}}{P_i^{\min}}$. Trade activity: $NbTrades_{\text{Client}}$ denotes the number of trades executed on behalf of clients, $NbTrades_{\text{Own}}$ is the number of trades executed for one's own account, $NbTrades_{\text{Liquidity}}$ is the number of trades executed as a liquidity provider (market maker), $NbTrades_{\text{RMO}}$ is the number of trades executed for Retail Market Organization (RMO), and $NbTrades_{\text{Others}}$ is the number of trades executed for other reasons. All variables are scaled to have comparable magnitudes. Z-score normalization transforms the variables to have a mean of 0 and a standard deviation of 1. Significance levels: ***: $p < 0.001$, **: $p < 0.01$, *: $p < 0.05$.

range 10 minutes before the circuit breaker activation (estimate = 1.066, p = 0.0206). This indicates that as the number of client trades increases, the relative range tends to widen, suggesting a possible acceleration in price movement as the circuit breaker approaches. However, the relationship becomes statistically insignificant after the market reopens, with an estimate of 0.619 and a p-value of 0.2292, indicating no notable effect on the price movement after reopening. Other variables, such as $NbTrades_{\text{Own}}$, $NbTrades_{\text{Liquidity}}$, $NbTrades_{\text{RMO}}$, and $NbTrades_{\text{Others}}$, do not show significant effects either before or after the circuit breaker event for pure-HFT traders.

In the case of mix-HFT traders, the results are more varied. The variable $NbTrades_{\text{Client}}$ shows a negative and statistically significant relationship before the circuit breaker (estimate = -1.831, p = 0.0122), indicating that an increase in client trades is associated with a decrease in the relative range in the pre-circuit breaker period. This suggests a stabilizing effect of client trades in the lead-up to the circuit breaker. However, this effect does not persist after the market reopens, as the estimate of 0.613 and $p = 0.1126$ indicate a lack of statistical significance. Other variables for the mix-HFT category, including those representing trades executed for own account and liquidity provision, do not show any significant impact on the relative range before or after the circuit breaker.

For non-HFT traders, the findings are more pronounced. The variable $NbTrades_{\text{Client}}$ shows a significant negative relationship before the circuit breaker (estimate = -1.102, $p = 0.0204$), suggesting that higher trading activity on behalf of clients leads to a reduction in the relative range, which could be interpreted as a mitigating effect on price volatility leading up to the circuit breaker. After reopening, the relationship becomes less clear, with an estimate of 0.335 and $p = 0.2342$, indicating no significant effect. Other variables, such as those representing own account trades and liquidity provision, show varying degrees of significance before and after the circuit breaker, with some results reaching significance, such as $NbTrades_{\text{Own}}$ (estimate = 0.860, $p = 0.0210$) before the circuit breaker.

The model's explanatory power, as indicated by the R^2 values, is 0.7415 before the circuit breaker, suggesting that the model explains a substantial portion of the variability in the relative range. After the circuit breaker, the R^2 drops to 0.4421, reflecting a decrease in explanatory power after the event.

These observations can be explained by the fact that most companies in this study have low-priced equities with high volatility usually avoided by HFTs (Hendershott *et al.*, 2011; Chung and Zhang, 2014; Budish *et al.*, 2015; Menkveld, 2013). High-frequency trading strategies typically rely on executing numerous small, rapid trades that capitalize on minor price movements. However, in markets with low liquidity, there may not be sufficient buy or sell orders to execute these trades efficiently without causing significant price impacts. Consequently, low liquidity increases the risk of being unable to enter or exit positions at desired prices, which is undesirable for HFTs. HFTs prefer stable and liquid environments where they can deploy high-speed algorithms to capture small price changes, and they tend to avoid markets with high volatility and low liquidity (Menkveld, 2013; Budish *et al.*, 2015).

In conclusion, the results suggest that the relationship between trade activity and price movement leading up to a circuit breaker varies across trader categories and time intervals. For pure-HFT and mix-HFT traders, the effects are mostly insignificant, while non-HFT traders exhibit more notable relationships with the relative range. These findings highlight the different dynamics at play among different types of traders, particularly in how they influence market conditions before and after a circuit breaker activation.

5 Leveraging Machine Learning to Enhance Market Stability

In this section, we explore the use of artificial intelligence (AI) tools to predict extreme market events via trading and order book activities that may trigger circuit breakers. At the core of AI is machine learning, a methodology where algorithms improve their ability to analyze and classify data through repeated exposure to historical datasets. This process, commonly referred to as training or learning, enables models to identify patterns and make accurate predictions on new data.

Machine learning-based studies typically follow two key phases. The first involves selecting relevant variables and optimizing models through a process of training and validation. A subset of the data is reserved for this purpose, allowing models to fine-tune their predictive capabilities. The second phase tests these optimized models on a separate dataset to assess their predictive performance.

Two principal methodologies dominate AI and machine learning applications: neural networks and decision trees. Neural networks operate in a supervised learning environment, processing inputs through multiple interconnected layers of nodes. These layers refine their outputs iteratively using backpropagation, adjusting weights to enhance prediction accuracy. Deep learning extends this approach, incorporating additional intermediate layers to process complex data patterns, which further improves prediction outcomes.

On the other hand, decision trees follow a more traditional approach by partitioning data sequentially. Each branch of the tree represents a decision split based on predictor variables, making this method intuitive and effective for many classification tasks. Both neural networks and decision trees exemplify AI's broader ability to recognize patterns and extract classification rules from historical data. Crucially, machine learning algorithms not only adapt to new information but also excel in generating predictions with minimal human intervention.

Since trades originate from activities within the order book, this section investigates the use of machine learning techniques to predict potential extreme market events, such as flash crashes or sharp bubbles, by examining trading patterns within the order book. To assess the contribution of each trading category-pure-HFT, mix-HFT, and non-HFT-the following metrics are analyzed: the average number of order modifications during the period; the order cancellation rate, defined as the number of canceled limit orders divided by the number of submitted limit orders; the buy/sell ratio, calculated as the number of bid orders divided by the number of ask orders during the period; the market/limit ratio of orders; the average number of trades associated with each order, representing the number of trades an order is involved in; the average number of trades executed by each trader category during the period; and the average trade volume for each trader category during the period. These metrics form the structure of the vector used for data processing by the machine learning algorithms.

5.1 Recurrent Neural Network (RNN) with ADAM Optimization Function

The dataset introduced in the previous section is both multivariate and temporally dependent, making it particularly challenging to analyze. To address this complexity, we employed a Recurrent Neural Network (RNN), which is well-suited for processing such data structures.

These models iterate repeatedly to optimize for the best approximation function between inputs and outputs. A key component of any neural network is its optimization function (Schmidhuber, 2015; LeCun *et al.*, 1998). For our model, we selected the ADAM (Adaptive Moment Estimation) optimizer (Kingma and Ba, 2014; Ruder, 2016). ADAM is a stochastic gradient-based optimization algorithm that improves upon traditional stochastic gradient descent by dynamically adjusting learning rates, handling noisy gradients effectively, and enabling faster convergence. These attributes make it especially suitable for training complex neural networks, such as ours, designed to predict potential incidents. During each iteration, ADAM operates by calculating gradients, updating moments, correcting biases, and finally updating parameters:

- **Moment Updates:**
  - Update the first moment (mean of the gradients):

$$m_t = \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g_t$$

  - Update the second moment (mean of the squared gradients):

$$v_t = \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t^2$$

- **Bias Correction:**
  - Correct the bias for the first moment:

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t}$$

  - Correct the bias for the second moment:

$$\hat{v}_t = \frac{v_t}{1 - \beta_2^t}$$

- **Parameter Updates:**

$$\theta_t = \theta_{t-1} - \frac{\eta \cdot \hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon}$$

where η is the learning rate, β_1 is the exponential decay rate for the first moment (mean), β_2 is the exponential decay rate for the second moment (variance), ϵ is a numerical stability term, m_t is the first moment estimate (initialized to 0), v_t is the second moment estimate (initialized to 0), and t is the iteration counter (initialized to 0).

To properly configure an RNN, it's crucial to carefully tune several hyperparameters: the learning rate (`lr`), the numerical stability term (ϵ), and the number of layers, to ensure optimal performance and avoid common issues like overfitting or underfitting. The learning rate controls how much the weights of the RNN are updated during each iteration of training. It is one of the most important hyperparameters because it directly affects the convergence of the model. If the learning rate is too high, the model may fail to converge or may oscillate around the optimal weight values, leading to poor performance. If the learning rate is too low, training can be slow and may get stuck in local minima, resulting in suboptimal results and increased training time. Additionally, too few layers might not have the capacity to capture complex temporal patterns in the data, leading to underfitting. On the other hand, too many layers (deep RNNs) can lead to overfitting, where the model memorizes the training data and fails to generalize well to unseen data. Deep RNNs can also suffer from vanishing gradients (when gradients become too small to update weights properly) or exploding gradients (when gradients become excessively large, destabilizing the training process).

Figure 4 shows that increasing the number of layers initially improves the validation loss but eventually leads to deterioration in performance. To counteract this, we implemented an early stopping mechanism that halts training when the validation loss increases for a specified number of iterations, which we set to 10 after conducting several tests. The learning rate (lr) also plays a significant role in convergence: the smaller the value, the faster the function converges, though it introduces more noise during training. Based on our experiments, we selected $lr = 0.0001$ as the optimal value. The term ϵ is used to prevent division by zero during training. If ϵ is too small, the model may become unstable, while if it is too large, it may diminish the effect of the corrected gradients. We chose $\epsilon = 10^{-8}$, as it provided a reasonable balance.

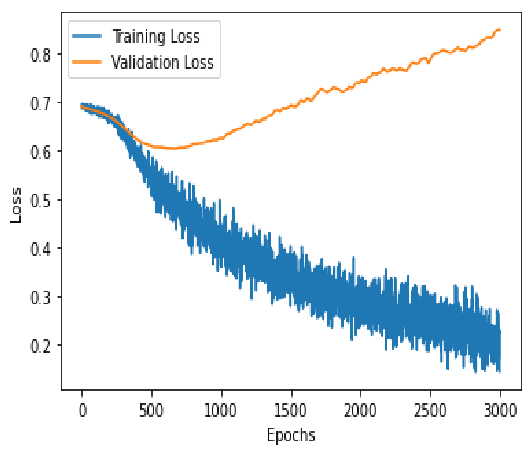

(a) $lr = 0.00001$, $\epsilon = 10^{-6}$, $Nb_{layer} = 3000$

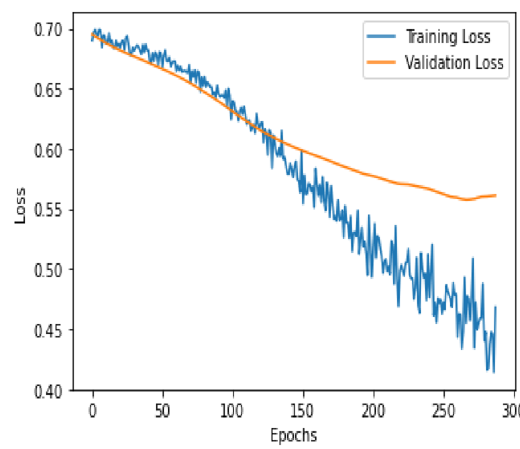

(b) $lr = 0.00001$, $\epsilon = 10^{-6}$, $Nb_{layer} = 300$

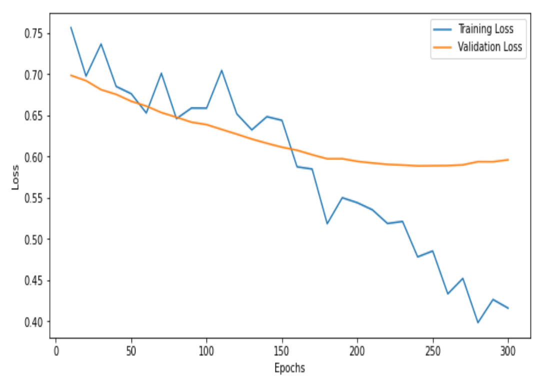

(c) $lr = 0.001$, $e = 10^{-5}$, $Nb_{layer} = 300$

Figure 4: The performance of the neural network as a function of the number of the training layers.

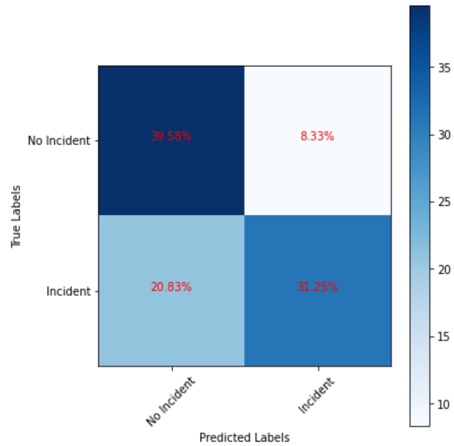

Figure 5: **Confusion matrix RNN algorithm.**

Using the values from the confusion matrix presented in the figure 5, the following metrics are calculated:

- **Accuracy**:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{Total}} = \frac{39.58 + 31.25}{39.58 + 8.33 + 20.83 + 31.25} = 70.83\%$$

  This indicates that the model correctly predicts 70.83% of the incidents.

- **Precision (for Circuit Breaker)**:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} = \frac{39.58}{39.58 + 20.83} = 65.51\%$$

  This means that when the model predicts a circuit breaker, it is correct 65.51% of the time.

- **Recall (for Circuit Breaker)**:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} = \frac{39.58}{39.58 + 8.33} = 82.59\%$$

  This suggests that the model identifies 82.59% of actual circuit breaker incidents.

- **F1 Score**:

$$\text{F1} = 2 \times \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} = 2 \times \frac{0.6551 \cdot 0.8259}{0.6551 + 0.8259} = 73.16\%$$

  This represents the harmonic mean of precision and recall.

where TP is the number of true positives, FP is the number of false positives, TN is the number of true negatives and FN is the number of false negatives.

The model demonstrates a high recall rate (82.59%), signifying its effectiveness in identifying actual circuit breaker events. However, its moderate precision (65.51%) suggests the presence of some false alarms, where the model predicts a circuit breaker that does not occur. The relatively high false positive rate (20.83%) indicates a tendency to overestimate the likelihood of circuit breakers. In financial markets, avoiding false negatives (missed circuit breakers) is often more critical than minimizing false positives (false alarms). Thus, the model's low false negative rate (8.33%) is a particularly favorable attribute.

To further refine the analysis, we employ a RNN algorithm to identify the features most influential in triggering circuit breakers. Figure 6 illustrates the individual order book activities across different trading categories, highlighting their respective impacts.

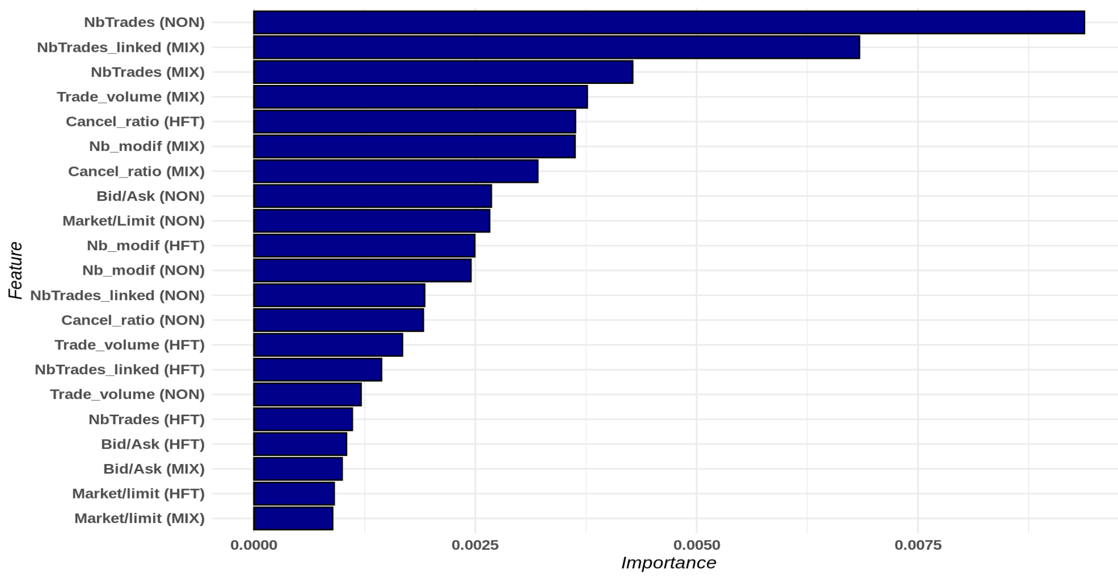

Figure 6: **The importance of features of RNN algorithm affecting predicted market extreme events leading to the circuit breaker triggering. Impact of order book and trading activity of each category of traders on the extreme events.**

Nb modif represents the average number of modified orders during the period. *Cancel ratio* refers to the ratio of canceled limit orders to submitted limit orders. *Buy/sell* indicates the ratio of bid orders to ask orders, while *Market/limit* shows the ratio of market orders to limit orders during the same period. *Order-to-trades* measures the average number of trades associated with each order, reflecting the number of trades an order is involved in. *Trade volume* represents the average trade volume for each trader category, and *Nb trades* denotes the average number of trades executed by each trader category during the period.

5.2 Random Forest

Another effective method for this type of data is the "Random Forest" algorithm. Random Forest is a supervised ensemble learning algorithm that combines multiple simpler models, typically decision trees, to achieve more robust and accurate performance. The algorithm works by training several decision trees, each built on a different subset of the training dataset (bootstrap sampling). Additionally, each tree is trained using a random subset of features, which improves model diversity and reduces correlation among the trees.

Once training is completed, the algorithm makes predictions in two ways:

- **For regression problems:** Random Forest averages the predictions of all the trees.
- **For classification problems:** Random Forest uses majority voting among the predictions of the trees.

The key parameters to consider include the number of trees in the "forest," the maximum depth of the individual trees, and the minimum number of samples required to split a node:

- **Number of trees:** Increasing this number often improves performance but also increases computational time and memory usage.
- **Maximum depth:** Limiting tree depth helps prevent overfitting, especially with noisy datasets.
- **Minimum number of samples required to split a node:** Lowering this value allows for more complex trees.

We used the `scikit-learn` library to implement our Random Forest. This library's built-in version is far more efficient and comprehensive than manually coding a Random Forest algorithm.

We chose to create a forest of 500 trees with a maximum depth of 20 and a minimum of 10 samples required to split a node. These parameters were selected empirically. Since our dataset is extensive, highly detailed, and noisy, overly complex models led to divergence in results, while overly simplistic models were less accurate.

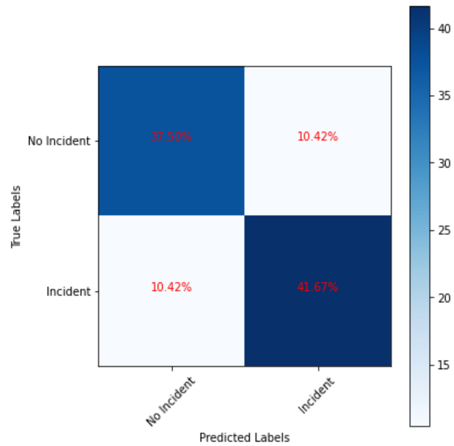

Figure 7: **Confusion matrix random forest algorithm.**

- The overall **accuracy** of the model is calculated as:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} = \frac{41.67 + 37.50}{37.50 + 10.42 + 10.42 + 41.67} = 79.17\%$$

  This indicates that the model correctly predicts about 79.17% of the cases.

- **Precision (Positive Predictive Value)** evaluates the reliability of positive predictions:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} = \frac{41.67}{41.67 + 10.42} = 80.00\%$$

  This means that 80% of the predicted circuit breakers are correct.

- **Recall (Sensitivity or True Positive Rate)** assesses the model's ability to identify actual circuit breakers:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} = \frac{41.67}{41.67 + 10.42} = 80.00\%$$

  The model detects 80% of actual circuit breaker events.

- The proportion of normal conditions misclassified as extreme events (**False Positive Rate**, FPR) is:

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} = \frac{10.42}{10.42 + 37.50} = 21.74\%$$

  This suggests a manageable level of overestimation of risk.

- The proportion of missed circuit breakers **(False Negative Rate, FNR)** is:

$$\text{FNR} = \frac{\text{FN}}{\text{FN} + \text{TP}} = \frac{10.42}{10.42 + 41.67} = 20.00\%$$

  This reflects a relatively low rate of missed events, which is crucial in financial markets.

The Random Forest model exhibits strong predictive performance, with balanced precision and recall rates both at 80%. This suggests that the model is effective in identifying extreme events leading to circuit breaker triggers, which is crucial for risk management in financial markets. The accuracy of 79.17% indicates that the model correctly classifies most of the market events, distinguishing between actual circuit breaker events and normal trading conditions. However, the presence of some false positives (FPR: 21.74%) indicates that the model tends to overestimate the likelihood of extreme events, which could result in unnecessary market interventions or a higher risk aversion.

Despite this, the model's ability to correctly identify 80% of the true circuit breaker events (Recall: 80.00%) is a critical strength, especially in environments where false negatives—missed circuit breakers—pose a higher risk. The model's relatively low false negative rate (FNR: 20.00%) demonstrates that it is capable of capturing the majority of extreme market movements, which could be vital for preventing financial crises or excessive volatility.

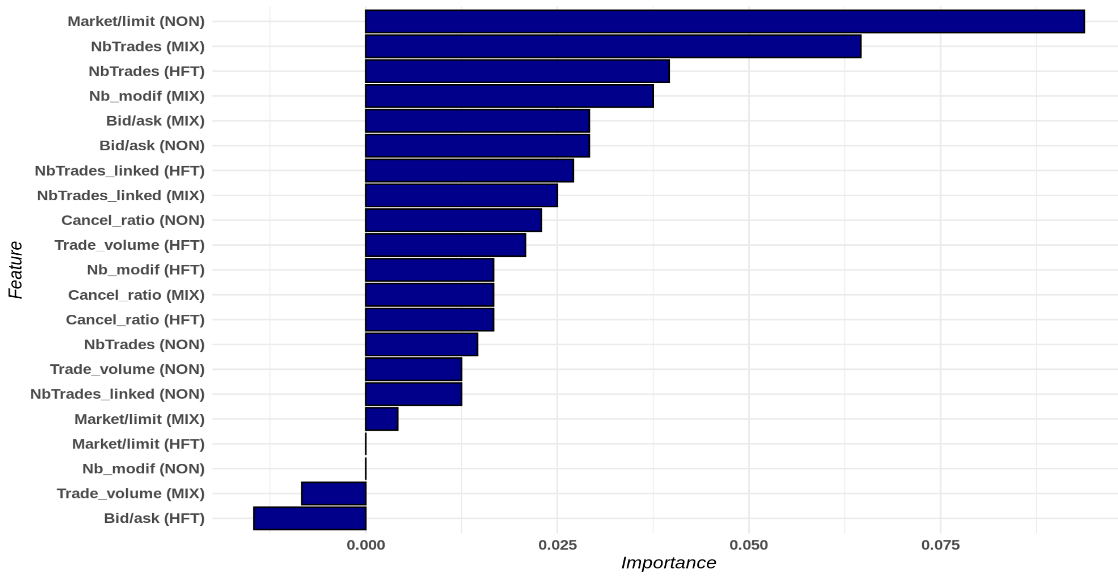

Figure 8: **Impact of order book activity of each category of traders relying on random forest algorithm to predict circuit breakers.**

Nb modif represents the average number of modified orders during the period. *Cancel ratio* refers to the ratio of canceled limit orders to submitted limit orders. *Buy/sell* indicates the ratio of bid orders to ask orders, while *Market/limit* shows the ratio of market orders to limit orders during the same period. *Order-to-trades* measures the average number of trades associated with each order, reflecting the number of trades an order is involved in. *Trade volume* represents the average trade volume for each trader category, and *Nb trades* denotes the average number of trades executed by each trader category during the period. Features with the negative impact are not informative for the prediction task.

6 Conclusion

The results of this paper reveal that non-HFTs are the most significant contributors to price instability leading to circuit breaker activations, particularly during downward trends. Surprisingly, HFTs, often criticized for destabilizing markets, demonstrate stabilizing behaviors during circuit breakers, especially in post-halt scenarios.

We also demonstrate that machine learning algorithms can be effectively utilized to predict events leading to circuit breaker activations by analyzing trading and order book activities of traders. The machine learning models explored in this study achieve high accuracy in predicting circuit breaker triggers, with RNNs achieving a recall rate of 82.59%, highlighting their effectiveness in identifying extreme events. However, both RNNs and random forests exhibit moderate precision, reflecting the challenge of minimizing false positives in volatile markets.

In the context of financial market regulation, where the ability to predict and mitigate the risks of flash crashes or bubbles is vital, the performance of the neural network and random forest models is promising. While further tuning and optimization of the algorithms may help reduce the false positive rate, their current performance strikes an effective balance between minimizing missed events (false negatives) and avoiding unnecessary alarms (false positives). This balance makes them useful tools for predicting circuit breaker triggers and supporting timely, data-driven interventions to stabilize market conditions.

Moreover, the models' robustness in detecting extreme events could be further enhanced by incorporating additional market features or refining the underlying algorithms. Future research could explore the integration of more advanced machine learning techniques, such as deep learning or ensemble models, to improve their ability to handle even more complex market dynamics.

Including additional features, such as order book depth or a global market sentiment index, could provide a more comprehensive understanding of market dynamics and trader behavior, potentially enhancing the predictive power of machine learning models.

References

Boonpan, S. and Sarakorn, W. (2025). Deep neural network model enhanced with data preparation for the directional predictability of multi-stock returns. *Journal of Open Innovation: Technology, Market, and Complexity*, **11**, 100438.

Brennan, M. J. (1986). A theory of price limits in futures markets. *Journal of Financial Economics*, **16**, 213–233.

Budish, E., Cramton, P. and Shim, J. (2015). The high-frequency trading 'arms race': Too fast to fail. *The Journal of Legal Studies*, **44** (2), 333–357.

Chen, H., Petukhov, A., Wang, J. and Xing, H. (2024). The dark side of circuit breakers. *The Journal of Finance*, **79**, 1405–1455.

Chowdhry, B. and Nanda, V. (1998). Leverage and market stability: The role of margin rules and price limits. *Journal of Business*, **71**, 179–210.

Chung, K.-H. and Zhang, H. (2014). High-frequency trading, stock volatility, and price discovery. *Journal of Financial Markets*, **18**, 1–27.

ESMA (2015). Esma, strategic orientation 2016-2020.

Gao, K., Vytelingum, P., Weston, S., Luk, W. and Guo, C. (2024). High-frequency financial market simulation and flash crash scenarios analysis: An agent-based modelling approach. *JASSS*, **27**.

Goldstein and Kavajecz (2004). Trading strategies during circuit breakers and extreme market movements. *Journal of Financial Markets*, **7**, 301–333.

Greenwald and Stein (1991). transactional risk, market crashes, and the role of circuit breakers. *The Journal of Business*.

GROSSMAN, S. J. and MILLER, M. H. (1988). Liquidity and market structure. *The Journal of Finance*, **43**, 617–633.

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

Jian, Z., Zhu, Z., Zhou, J. and Wu, S. (2020). Intraday price jumps, market liquidity, and the magnet effect of circuit breakers. *International Review of Economics and Finance*, **70**, 168–186.

Kim, K. A. and Rhee, S. G. (1997). Price limit performance: Evidence from the tokyo stock exchange. *The Journal of Finance*, **52**, 885–901.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kolte, A., Roy, J. K. and Vasa, L. (2023). The impact of unpredictable resource prices and equity volatility in advanced and emerging economies: An econometric and machine learning approach. *Resources Policy*, **80**, 103216.

Leal, S. J. and Napoletano, M. (2019). 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*, **157**, 15–41.

LeCun, Y., Bottou, L., Bengio, Y. and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, **86** (11), 2278–2324.

Lee, J. and Schu, L. (2022). Regulation of algorithmic trading: Frameworks for human supervision and direct market interventions *.

Li, Z., Hou, K. and Zhang, C. (2021). The impacts of circuit breakers on china's stock market. *Pacific-Basin Finance Journal*, **68**, 101343.

MENKVELD, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, **16** (4), 712–740.

RIORDAN, R. and STORKENMAIER, A. (2012). Latency, liquidity and price discovery. *Journal of Financial Markets*, **15**, 416–437.

RUDER, S. (2016). An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.

SCHMIDHUBER, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, **61**, 85–117.

SIFAT, I. M. and MOHAMAD, A. (2020). A survey on the magnet effect of circuit breakers in financial markets. *International Review of Economics and Finance*, **69**, 138–151.

SUBRAHMANYAM, A. (1994). Circuit breakers and market volatility: A theoretical perspective. *Journal of Finance*, **49**, 237–254.

— (2013). Algorithmic trading, the flash crash, and coordinated circuit breakers. *Borsa Istanbul Review*, **13**, 4–9.

WANG, S. S., XU, K. and ZHANG, H. (2019). A microstructure study of circuit breakers in the chinese stock markets. *Pacific Basin Finance Journal*, **57**.

WANG, X., KIM, M. H. and SUARDI, S. (2022). Herding and china's market-wide circuit breaker. *Journal of Banking and Finance*, **141**.

WONG, K. M., KONG, X. W. and LI, M. (2020). The magnet effect of circuit breakers and its interactions with price limits. *Pacific Basin Finance Journal*, **61**.

YUAN, F. (2024). Research on time-series financial data prediction and analysis based on deep recurrent neural network. *Applied and Computational Engineering*, **69**, 140–146.