

Article

When Will the Next Shock Happen? A Dynamic Framework for Event Probability Estimation

Konstantinos Pantelidis , Ioannis Karakostas  and Odysseas Pavlatos 

Department of Accounting and Finance, University of Macedonia, 54636 Thessaloniki, Greece; karakostas@uom.edu.gr (I.K.); opavlatos@uom.edu.gr (O.P.)

* Correspondence: kpantelidis@uom.edu.gr

Abstract

Extreme movements in financial time series pose challenges for risk management and forecasting, particularly when their timing is irregular and difficult to anticipate. This study aims to develop a probabilistic framework for detecting and predicting such events using daily Bitcoin returns as a case study. We first identify extreme positive and negative return events using the Isolation Forest algorithm and estimate their empirical recurrence patterns using a dynamic frequency table to derive baseline parametric probabilities. A 7-day Hawkes excitation kernel is then applied to capture short-run self-exciting dynamics, and both components are integrated using logistic regression to produce real-time probability forecasts. The results show that positive events occur more frequently than negative ones and that prediction accuracy improves over time: Brier scores, which measure the accuracy of probabilistic predictions, decrease as additional event data accumulate, and log loss values exhibit a consistent downward trend. Overall, by combining anomaly detection, empirical inter-arrival estimation, and excitation dynamics into a unified structure, the proposed framework offers a transparent and adaptable tool for forecasting extreme events in the financial market.

Keywords: event detection; Hawkes process; probability estimation; logistic regression

JEL Classification: C22; C58

1. Introduction

Financial time series experience both favorable and adverse events, leading to certain days with exceptionally high or low returns that stand out from the typical market behavior. This raises the question of whether we can identify these anomalous days as discrete events and predict their likelihood of recurrence by analyzing the intervals between them. This study explores such an approach by identifying extraordinary days in Bitcoin's return data and computing a baseline probability for positive and negative events. We then applied a 7-day Hawkes excitation kernel to refine these estimates, ultimately combining both measures through logistic regression to generate real-time probability forecasts for positive and negative market events.

Event study analysis is applied in various fields of science, including finance, economics, engineering, seismology, and epidemiology [1–4]. An event could indicate positive news, but it could also signify a catastrophe that could be prevented or at least lead to better preparation for affected parties. While the definition of an event is not always the same, the concept remains constant; all attempts focus on understanding why an event



Academic Editor: Abdullahi Dahir Ahmed

Received: 13 October 2025

Revised: 21 December 2025

Accepted: 30 January 2026

Published: 2 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

occurs, which could be depicted as a break in a time series, its potential impact, and how a future event may be predicted.

Although detecting events on time series is an ongoing and extensively studied topic, with various methods used to achieve this with varying levels of efficiency [5], investigating alternate approaches to calculating the probabilities that these events could appear at a given time could be a valuable addition to the topic. Although focusing on the volatility of asset returns, hedge fund managers and traders who want to mitigate risk in their investments could benefit from tools that anticipate extremely volatile price events based on their timely recurrence. Moreover, policymakers, exchanges, and institutions responsible for monitoring systemic risks could benefit from forward-looking signals that indicate when markets are entering periods of heightened instability. Existing approaches, such as GARCH-type volatility models, Extreme Value Theory, and classical Hawkes processes, capture important aspects of tail behavior or clustering; however, they generally do not integrate anomaly detection, inter-arrival patterns, and excitation dynamics into a unified, real-time forecasting framework. This gap motivates the need for models that combine event identification with probability estimation, offering a clearer operational value for users who must anticipate rare but consequential market movements. Subsequently, this study aims to address this issue by proposing an event detection framework that focuses on predicting the probability of an event occurring, specifically in the context of financial time-series data.

To clarify the contribution from the outset, the proposed framework comprises four stages. First, a machine learning anomaly detection method identifies extreme return events without relying on rigid statistical thresholds. Second, the empirical distribution of inter-arrival times provides a dynamic baseline for the long-run event frequency. Third, short-term clustering effects are incorporated using a Hawkes-inspired excitation mechanism that raises event probability following recent shocks. Finally, these components were combined using logistic regression to generate real-time probability forecasts. This outline previews how the study brings together distinct strands of literature into a unified, forward-looking event-probability model.

The remainder of this paper is organized as follows: Section 2 presents a literature review, followed by the methodology in Section 3. The results are presented in Section 4 and discussed in Section 5, before concluding in Section 6.

2. Literature Review

Traditionally, the analysis of financial events begins with the event study methodology proposed by MacKinlay (1997) [6]. His framework assesses how specific events affect stock prices by comparing the actual and counterfactual returns within a structured event window. Although foundational, this approach is inherently backward-looking; it measures how markets react after an event occurs without offering insight into when such events might emerge. This limitation recurs across much of the literature and motivates more forward-looking modelling efforts in the future.

To address changes that unfold over time, researchers have focused on identifying such structural breaks. Bayesian change-point methods (Barry & Hartigan, 1993 [7]; Fearnhead, 2006 [8]) update beliefs about regime shifts as new observations arrive, treating the series as segments governed by different dynamics. Similarly, Hamilton's (1989) Markov regime-switching model [9] captures the transitions among latent economic states. These methods detect underlying shifts but typically require a full sample and impose substantial computational burdens, making them difficult to implement in real time. Most importantly, they do not provide a direct probabilistic estimate of the timing of future events.

Volatility models add another dimension to this analysis. Autoregressive conditional heteroskedasticity (ARCH) (Engle, 1982) [10] and generalized ARCH (Bollerslev, 1986) [11] demonstrate volatility clusters, while exponential GARCH (Nelson, 1991) [12] and threshold GARCH (Zakoian, 1994) [13] introduce asymmetric responses to negative shocks [14]. These frameworks successfully quantify risk dynamics and the intensity of market stress but remain focused on magnitude rather than event recurrence, offering limited guidance on the arrival process of disruptive episodes.

Extreme Value Theory [15] similarly characterizes the tail behavior of returns, informing us about the severity and frequency of extreme outcomes. However, this theory does not address temporal dependence or the likelihood of future events occurring. It describes the magnitude of an extreme event, not when it is likely to occur.

Duration-based and hazard models bring the analysis closer to modelling the timing of the events. Cox's proportional hazards framework (1972) [16] and financial analogues such as Shumway's (2001) bankruptcy model [17] estimate the probability of an event occurring within a future interval. However, they rely on parametric hazard functions, typically exponential or Weibull functions, which may not capture the complex, nonlinear, and clustered nature of these events. Their structure lacks a mechanism for short-term contagion, which is a well-documented empirical regularity in financial data series.

Stochastic models incorporating discontinuities, such as Merton's (1976) jump-diffusion [18] and Bates' (1996) stochastic volatility model with jumps [19], explicitly represent abrupt changes in prices. They distinguish between continuous and jump components but do not specify a data-driven or self-adjusting process that governs the arrival of the jumps.

A significant advance in explicitly modelling event recurrence appears in Hawkes' (1971) self-exciting process [20]. Hawkes models capture the empirical fact that events often cluster: one event temporarily elevates the probability of another, which then decays over time. Their structure combines the baseline intensity with an excitation mechanism, offering a natural way to model contagion. However, most implementations assume a constant or slowly varying baseline intensity, limiting their ability to capture long-term trends or structural shifts known to exist in financial instability.

Building on this foundation, a growing body of literature applies Hawkes processes to cryptocurrency markets, where event clustering and feedback effects are particularly pronounced. Phillips and Gorse (2018) [21] employed a multivariate Hawkes framework to study the mutual excitation between cryptocurrency returns and social media topics, treating both price movements and sentiment indicators as interacting event streams. Luo et al. (2022) [22] model Bitcoin block arrivals jointly with price jumps, using Hawkes processes as a generative mechanism to analyze self- and cross-excitation between blockchain activity and price discontinuities. More recent contributions extend Hawkes-based modeling to high-frequency environments, such as the study of Cestari et al. (2024) [23], who apply Hawkes processes to limit order book data to forecast microstructure event timing, while Yang et al. (2025) [24] use multivariate Hawkes models to study jump contagion between cryptocurrencies, equity markets, and market sentiment indicators.

Crisis prediction studies provide a different perspective. Kaminsky and Reinhart (1999) [25] identify threshold signals for banking crises, while Borio and Drehmann (2009) [26] show that cumulative indicators enhance predictive performance. These approaches produce valuable early warning tools but depend heavily on macroeconomic aggregates rather than endogenous event dynamics within the financial time series.

Finally, machine learning approaches, such as Isolation Forests [27] and Long Short-Term Memory networks (Malhotra et al., 2015) [28], detect anomalies and latent events [29] using data-driven pattern recognition. Their flexibility is substantial; however, they lack a clear and interpretable probabilistic structure for event timing and often treat anomalies as isolated observations rather than components of a self-reinforcing or duration-dependent mechanism.

Related studies in other scientific domains also illustrate the usefulness of state-dependent event activation. For instance, Zhu et al. (2024) [30] developed a dynamic event-triggered mechanism for state estimation in robotics, in which events were generated only when the underlying conditions crossed the adaptive thresholds. In a similar spirit, Gao et al. (2024) [31] proposed a dynamic event-triggered monitoring and control architecture for discrete-time Markovian jump systems and demonstrated how transmissions can be activated endogenously by regime-dependent system states. Additional evidence that event occurrences can be governed by system-specific dynamics rather than fixed schedules comes from Phuntsho and Gonsalves (2024) [32], who examine event-based payment structures within resource-constrained project scheduling using approximate dynamic programming. Although situated in a different application domain, their formulation reinforces the broader methodological insight that event timing can emerge endogenously within a dynamic system, a principle that aligns with the hybrid duration-and-self-exciting structure pursued in this framework.

Collectively, the literature excels at measuring reactions to events, identifying structural breaks, modelling volatility dynamics, and detecting anomalies. However, none of these studies combine long-term duration dependence with short-term self-excitation in a unified, interpretable framework capable of estimating the probability of future events in real time. The lack of a model that combines (i) gradual baseline changes, (ii) rapid contagion effects, and (iii) nonparametric, data-driven timing dynamics inspired the creation of our proposed approach. This approach integrates duration modeling with Hawkes-type mechanisms to generate forward-looking event probabilities based on both long-term and short-term clustering.

3. Materials and Methods

3.1. General Framework

This section presents the general framework of the proposed methodology. Cryptocurrency market data were retrieved from the CoinGecko public API (Gecko Labs Pte. Ltd., Singapore) [33] and analyzed using Python (Python Software Foundation, Wilmington, DE, USA) [34], with standard data analysis libraries such as NumPy [35] and pandas [36]. Initially, we need a time series to be analyzed to define a number of n events E at time t . Assume that we obtain our time-series data, which may be prices and their returns, dynamically, with constant updates. Therefore, in each step at time t , we decide whether event i is observed compared to a set of observations w from $t-j$, where j denotes the lag (in discrete time steps) between the current time t and the start of the window. Suppose an observation contains an event, then $E(i, t) = 1$, else $E(i, t) = 0$. Next, we present a series of tables in which we insert the results. First, we have a table of $t \times n$ observations in which we update the result of the presence of an event in each present step, leading to a histogram that records how long ago each occurrence occurred. Second, a table of $t \times n$ observations that includes the time elapsed m from one event to the other, with a minimum value of 1 for each observation is created. If an event detection occurs, that is, $E(i, t) = 1$, then $m_{t+1} = 1$ and $m_t = m_{t-1} + 1$; otherwise, $m_t = m_{t-1} + 1$. Third, we have n tables for each i event, which contain the frequency of the time elapsed before the event occurred. These tables contain the value of m_t when $E(i, t) = 1$ and are updated only when such an event occurs. Fourth,

for each frequency table, we calculated the empirical distribution and compared it with a set of theoretical distributions to find the best fit.

$$P(E \leq m)_{t,TVD} = \int_0^m f(e)_{t,DDF} de \quad (1)$$

Fifth, we calculate the cumulative probability function of the selected theoretical distribution, where the probability of an event at the minimum elapsed time is obtained by a dynamic distribution fitting (DDF) in Equation (1). The DDF could be, for example, the empirical distribution that is recalculated every time an event occurs or a theoretical distribution applied to the frequency table. Finally, we calculate the probability under the theoretical distribution at time t that an event would occur, given the time elapsed since the last event.

$$Exc(t, \beta) = \sum_{i:t_i < t} \exp(-\beta(t - t_i)) \quad (2)$$

Next, for each day, we define the exponential excitation function as shown in Equation (2), where t_i represents the date an event occurred, and t is the current date; therefore, we have the difference in days after an event occurred, and β is the decay factor. Each past event contributes an exponentially decaying signal; therefore, events far in the past contribute nearly zero excitation.

$$\text{logit } p_t^{Exc} = \gamma_0 + \gamma_1 Exc_t \quad (3)$$

Next, we fit the logistic regression model of the following form, as shown in Equation (3), to obtain a transition probability, where γ_0 is the intercept and γ_1 is the coefficient of the excitation term. The logistic regression model can be calculated using an expanding window approach to include all past events up to time t . Finally, we obtain a combined probability on the log-odds scale, as in Equation (4).

$$\text{logit } p_t^{Combined} = (1 - w) \cdot \text{logit } p_t^{Base_0} + w \cdot \text{logit } p_t^{Exc} \quad (4)$$

In (4), w is the weight assigned to the excitation component probability.

$$p_t^{Combined} = \text{Logistic}(\text{logit } p_t^{Combined}) = \frac{1}{1 + e^{-(\text{logit } p_t^{Combined})}} \quad (5)$$

Subsequently, the combined probability is derived from the inverse logit function, as shown in Equation (5).

3.2. Empirical Application

In this section, we present the empirical application of the proposed methodology.

In Figure 1, we present a flowchart of the applied methodology, which describes the methodology in Section 3.1 and how it will be implemented. The experiment was conducted in two stages: in the first stage, we showed how we collect the data and make the necessary calculations for outlier detection; in the second stage, we computed the probabilities and relevant performance metrics.




Figure 1. Methodology flowchart.

3.2.1. Outlier Detection

The first stage of the methodology involves the outlier detection mechanism.

In Algorithm 1, we present the outlier detection procedure for the experiments. Initially, we start by collecting data from Coingecko.com [33] on the daily price returns of Bitcoin from 1 May 2013, to 11 November 2025.

$$r_{BTC,t} = \ln(\text{Price}_{BTC,t}) - \ln(\text{Price}_{BTC,t-1}) \quad (6)$$

We calculated the natural log price returns as shown in Equation (6). Subsequently, we implemented a real-time simulation for the entire period of T observations, obtaining a set of values for a rolling window, $W_t = \{r_{BTC,t-29}, r_{BTC,t-28}, \dots, r_{BTC,t}\}$, of 30 observations starting at $t = 30$, advancing one step at a time for the entire sample. In each window, we identified price return outliers using the Isolation Forest algorithm [27,37].

Algorithm 1: Rolling Window Isolation Forest for Bitcoin Price Anomaly Event Detection

Require: Historical price data stream P , rolling window size $w = 30$, contamination level c , estimators es

```

1: Initialize historical dataset  $D_h$  from storage if exists, else  $D_h \leftarrow \emptyset$ 
2: Fetch new price data  $P_{\text{new}}$  via API call to cryptocurrency exchange
3: Preprocess  $P_{\text{new}}$ :
4:   Compute log returns:  $r_t = \log(P_t) - \log(P_{t-1})$ 
5:   Filter to relevant time period:  $D_{\text{new}} \leftarrow P_{\text{new}}$  [1 May 2013 to 11 November 2025]
6:   Copy past event labels from  $D_h$  to  $D_{\text{new}}$  avoiding future data, if  $D_h \neq \emptyset$ 
7: for each time index  $t = |D_h| + 1$  to  $|D_{\text{new}}|$  do
8:   Extract rolling window:  $W_t \leftarrow \{r_{t-w+1}, \dots, r_t\}$ 
9:   Train Isolation Forest model:
10:    $\text{IF}_t \leftarrow \text{IsolationForest}(n\_estimators = es, contamination = c, random\_state = 123)$ 
11:    $\text{IF}_t.\text{fit}(W_t)$ 
12:   Predict anomaly label for latest observation:
13:    $a_t \leftarrow \text{IF}_t.\text{predict}(r_t)$ 
14:   Label event:
15:    $\text{Event}_t \leftarrow 1$  if  $a_t = -1$  (outlier), else 0
16: end for
17: Classify event direction:
18:   for each event index  $e$  where  $\text{Event}_e = 1$  do
19:     If  $r_e > 0$  then
20:        $\text{PositiveEvent}_e \leftarrow 1, \text{NegativeEvent}_e \leftarrow 0$ 
21:     else if  $r_e < 0$  then
22:        $\text{PositiveEvent}_e \leftarrow 0, \text{NegativeEvent}_e \leftarrow 1$ 
23:     else if  $r_e = 0$  then
24:        $\text{PositiveEvent}_e \leftarrow 0, \text{NegativeEvent}_e \leftarrow 0$ 
25:     end if
26:   end for
27: Persist and Update DataFrame with detected events to storage

```

Isolation Forest is an unsupervised machine learning algorithm that identifies anomalies by recursively partitioning the data space, such that anomalous observations are isolated with fewer splits than regular observations. In each rolling window, the model is trained on the returns and used to evaluate whether the most recent return is anomalous,

where such observation is classified as an outlier when the algorithm assigns it an anomaly label of -1 .

The parameters for the study include a contamination factor of c , a random seed of “123,” and es estimators. In this study, we present detailed results for $es = 2$ and $c = 0.01$, while a summary of the sensitivity analysis performance metric results is included in the Appendix A (for $es = \{2, 50, 100, 200\}$ and $c = \{0.01, 0.05\}$). The contamination factor sets the assumed proportion of outliers in the data, while the estimators are isolation trees whose individual anomaly scores are averaged to produce a final anomaly prediction.

We defined two types of events: positive events, where the last observation of the set was classified as an outlier and the returns were positive, and negative events, which occurred when the returns were negative. In this way, we ensured that the algorithm would not consider future values of the returns, as in real-time situations.

3.2.2. Probability Estimation

Whenever an event occurred, the two tables corresponding to the positive or negative event distribution were recalculated the next day to obtain the correct baseline probability value. We selected the best-fitting distribution to the empirical distributions, which could be used for modelling heavy and thin data. The identified theoretical distributions that might effectively capture heavy-tailed data include the exponential, exponential power, gamma, power law, log-normal, and normal distributions for thin-tailed data. Heavy-tailed distributions are favored because the normal distribution tends to cluster probabilities near the average, which can lead to an underestimation of the likelihood of extreme events. The optimal parameters for each distribution that were tested for fit were calculated using a maximum likelihood estimation approach in Python version 3.10.9 [34] by utilizing the Fitter and SciPy libraries [38,39]. The selection of the optimal distribution was determined by selecting the one with the lowest Akaike Information Criterion value, ensuring minimal expected loss functions even when the true data-generating distribution was absent from the set of potential models or the best approximation was highly parameterized [40], while balancing the bias-variance issues that arose when fitting data [41]. In Equation (2), we assign the β variable a value of 0.10, which represents a half-life of approximately 7 days for the excitation factor. Next, we applied a weight of 0.05 to Equation (3), and the remaining 95% was assigned to the base probability, which was derived from Equation (1). We conducted logistic regression using the L1 penalty function and an expanding-window approach, ensuring that at least 60 observations were included, and updated the regression at each time step to factor in historical data.

$$\text{Brier Score} = \frac{1}{N} \sum_{t=1}^n (f_t - o_t)^2 \quad (7)$$

The accuracy of the probabilistic forecasts is evaluated using the Brier score [42], which is defined as in Equation (7). where f is the forecasted probability of the event at time t , and o is the outcome, which takes the value of 1 if the event occurs, and 0 otherwise. Lower Brier scores indicate better probabilistic calibration and accuracy.

$$\text{Log loss} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)] \quad (8)$$

Next, the logistic loss formula [43] in Equation (8) obtains the probability p at index i for N observations and a y label, which in this case is 0 and 1, noting the existence of an event at the examined index. Log loss is a performance metric for classification models that measures the discrepancies between the predicted probabilities and label values. The

value of this metric ranges from 0 to infinity, where values close to zero indicate higher confidence in the predictions. We calculated the Brier and Log loss metrics daily by utilizing an expanding window approach, incorporating all available data up to each point in time.

3.3. Model Performance on Extreme Returns

To evaluate the model's performance in relation to extreme returns, we adopted the following approach. Initially, after obtaining results from the detection of positive and negative events and their probability estimates, we excluded all returns that fell below or above zero, filtering for dates where a combined probability was available. We then compute the absolute values of these returns and select the top 5% of both positive and negative returns during the sample period, which gives us the extreme returns for each case. Subsequently, we calculate the Pearson and Spearman correlations to determine if there is a relationship between the combined probability and extreme returns, in addition to assessing the Brier and Log Loss scores for the entire sample in each case.

4. Results

4.1. Positive Events

The summary statistics for positive events are presented in Table 1.

Table 1. Positive events descriptive statistics.

Variable	Mean	Standard Deviation	Median	Min	Max	Observations
Total returns	0	0.04	0	−0.43	0.29	4576
Positive Returns	0.02	0.03	0.02	0	0.29	2426
Positive Event Days Elapsed	41.34	34.97	32	1	207	4576
Baseline Probability	0.51	0.24	0.53	0	1	4508
Excitation Probability	0.02	0	0.02	0.01	0.04	4508
Combined Probability	0.48	0.23	0.48	0	1	4508

As shown in Table 1, the baseline probability is the dominant factor in the model, averaging 0.51 compared to a much smaller excitation probability of 0.02, which results in a combined probability of 0.48. The analysis is based on 4576 total return observations, with probabilities computed for 4508 days, indicating that positive market events occur, on average, every 41 days.

The results for positive events are shown in Figure 2. We had a total of 94 positive events, with an average event occurrence after the first event of 48 days. In most cases, the proposed methodology successfully captures positive return spikes. The positive events spanned from 8 July 2013, to 2 October 2025. From 9 July 2013, to 28 May 2016, the baseline probability was derived from the exponential power, exponential, and normal distributions. A significant shift occurred on 29 May 2016, when the power-law distribution became dominant and remained so until 4 April 2018, followed by a 7-day period in which the gamma distribution was selected, which then alternated with the exponential power distribution until 1 May 2020. From that date onward, the power-law distribution again became the permanently dominant source of the baseline probability. This shows that in the presence of more data, the estimated chance of an extremely positive jump increases as the distributions shift from more thin-tailed to more fat-tailed, treating extreme events as more probable than impossible. The excitation probability remained low, ranging from a maximum of 0.04 to 0.01.




Figure 2. Positive events: full dataset overview.

Regarding the combined probabilities, we observed 48 detected positive events when the combined probability exceeded the threshold of 50%. However, in the first few months, the probability values seem to have unnatural jumps due to the lack of data on events. The performance analysis shows a mean Brier score of 0.27, which was approximately 0.06 in early 2014, climbed sharply to approximately 0.18 by February 2014, and maintained this level until early 2016, when it underwent a second increase, peaking at approximately 0.31 before gradually declining to 0.28. Mirroring this trend, the Log Loss score followed a similar trajectory, rising sequentially to values of 1.00 and 1.68 during the corresponding periods before eventually receding to 0.96, with a mean of 1.05.

The summary statistics for positive tail events are presented in Table 2.

Table 2. Extreme positive returns probabilities.

Variable	Mean	Standard Deviation	Median	Min	Max	Observations	Jarque–Bera
Price returns	0.11	0.03	0.10	0.08	0.29	120	408.56 ***
Abs. Price returns	0.11	0.03	0.10	0.08	0.29	120	408.56 ***
Combined Probability	0.40	0.25	0.42	0.00	1.00	120	3.37

Note. The asterisks “***” denote significance $p < 0.01$, rejecting the null hypothesis of normally distributed data.

As we observe in the summary statistics, both price returns and absolute price returns have identical distributions, with a mean of 0.11 exhibiting non-normality. In contrast, the combined probability had a lower mean of 0.40 and was much closer to a normal distribution. In total, 41 positive events were detected out of 120 observations. Figure 3 shows the plots of the extreme returns and the respective probability values.




Figure 3. Isolated positive events and returns.

The initial visual analysis does not identify a systematic pattern connecting extreme returns to higher probability values. Therefore, we continue with the results of the correlation analysis presented in Table 3.

Table 3. Extreme positive returns correlation analysis.

Correlation Test	ρ	p
Pearson	−0.077	0.404
Spearman	−0.054	0.556

Note. The analysis was based on 120 observations. The tested variables were absolute returns and combined probabilities.

The results of the correlation analysis show that there is no significant relationship between the examined variables, as the correlation coefficients, while negative, are close to zero, with statistically insignificant probability values.

4.2. Negative Events


The summary statistics for negative events are presented in Table 4.

The summary statistics indicate that Negative Returns range from −0.43 to 0, with a mean of −0.02. The baseline probability had a mean of 0.40, contributing the most to the combined probability with a mean of 0.37. In contrast, the excitation probability was trivial, with a mean of 0.01 and no variation. The analysis covers 4576 return observations and a history of Negative Event Days Elapsed ranging from 1 to 392, with a mean negative event appearance every 117 days.

Table 4. Negative events descriptive statistics.

Variable	Mean	Standard Deviation	Median	Min	Max	Observations
Negative Returns	−0.02	0.03	−0.01	−0.43	0	2150
Negative Event Days Elapsed	116.82	94.81	89	1	392	4576
Baseline Probability	0.40	0.31	0.32	0	1	4356
Excitation Probability	0.01	0	0.01	0	0.01	4356
Combined Probability	0.37	0.3	0.28	0	1	4356

The results regarding negative events are presented in Figure 4, which shows a total of 28 negative events starting from 7 December 2013, and ending on 20 December 2024. The initial distribution, which serves as the source of the baseline probabilities, is derived from the exponential power, followed by the normal distribution spanning approximately thirteen months to April 2016. Subsequently, there is a phase of fluctuating distributions that includes exponential power, gamma and normal distributions, continuing until an exponential power distribution is established for a period of three years starting in July 2017. For the rest of the examined period, the normal, gamma, and exponential power distributions alternate in capturing the tail behavior of the negative returns, alternating from a more relaxed behavior in the early sample period, implying less extreme elapsed days probability, becoming more elevated as the dataset expands, and then reverting to intermittent periods of moderated tail risk, potentially indicating a shift in the underlying frequency of negative events. The excitation probability remains at low levels of approximately 0.006. The combined probabilities have a mean of approximately 37%, whereas nine of the 28 negative events occurred when the probability was over 50%. The Brier score rises to approximately 0.43 in January 2015, then decreases until it stabilizes at approximately 0.20 from August 2018 onward, with a mean value of 0.20 for the entire period. Similarly, the log loss values rise up to 5.60 until they slowly decline to 1.09 at the end of the examined period, with a mean value of 1.67.

**Figure 4.** Negative events: full dataset overview.

The summary statistics for negative tail events are presented in Table 5.

Table 5. Extreme negative return probabilities.

Variable	Mean	Standard Deviation	Median	Min	Max	Observations	Jarque–Bera
Price returns	−0.12	0.05	−0.10	−0.43	−0.08	104	−0.12 ***
Abs. Price returns	0.12	0.05	0.10	0.08	0.43	104	0.12 ***
Combined Probability	0.35	0.30	0.32	0.00	1.00	104	0.35 **

Note. The asterisks “***” and “**” denote significance $p < 0.05$ and $p < 0.01$, respectively, rejecting the null hypothesis of normally distributed data.

The summary statistics show that the price returns for negative tail events have a mean of −0.12 and range from −0.43 to −0.08. The absolute price returns have a mean of 0.12 and range from 0.08 to 0.43. The combined probability showed a variation with a mean of 0.35 and a full range from 0.00 to 1.00. The Jarque–Bera test statistics indicate that the null hypothesis of normality is rejected at the 1% significance level for both return series and at the 5% level for the Combined Probability. In total, 17 negative events were detected out of 104 observations. A preliminary visual assessment of the probability and return values, as presented in Figure 5, does not support a clear correlation between the magnitude of returns and their corresponding probabilities.



Figure 5. Isolated negative events and returns.

Next, we present the results of the correlation analysis for negative extreme returns.

As shown in Table 6, there is no significant correlation between extreme negative returns and combined probabilities, with small negative coefficients and insignificant probability values.

Table 6. Extreme negative returns correlation analysis.

Correlation Test	ρ	p
Pearson	−0.044	0.658
Spearman	−0.092	0.354

Note. The analysis was based on 104 observations. The tested variables were absolute returns and combined probabilities.

5. Discussion

5.1. Key Findings

In a streaming data simulation, the suggested methodology effectively captures both positive and negative outliers in Bitcoin returns; nevertheless, the number of events is small in either scenario. Different baseline probabilities result from positive events occurring more frequently than negative events. The excitation component has a slight influence on the combined probabilities of both positive and negative events. Finally, as the event data grew, the Brier score in both scenarios decreased before stabilizing at lower values.

5.2. Interpretation of Results

The first observation is that the methodology successfully captures positive and negative return outliers in simulated streaming daily data by utilizing a 30-day window. Next, the difference in the event distributions indicates different behaviors in upward and downward moves, as the former are more frequent, whereas the latter are sparser. In addition, the excitation factor slightly adjusts the combined probabilities, providing a baseline calibration without inflating or deflating the baseline probabilities.

Regarding the Brier score for the positive and negative event data, it was initially explained by the low number of event observations necessary for calculating the baseline probabilities and partially by the excitation-adjusted probabilities that declined over time. Additionally, the rising and then falling Brier score is evidence of successful calibration over time with the addition of more events in the data, rather than failure. The log loss values have a similar pattern, as they tend to fall as more data become available, which should be interpreted as a sign of successful calibration over time with the addition of more events.

The correlation analysis revealed no statistically significant relationship between the model's probabilities and the tails of extreme positive and negative returns. This suggests that, in its current form, the model does not effectively capture the drivers or forecast the intensity of these tail events, limiting its utility for assessing extreme market moves.

5.3. Comparison with Literature

The framework proposed in this study links a multitude of methodologies, including event study analysis, volatility modeling, and hazard-based duration methodologies. Classical event study designs, such as the one introduced by MacKinlay [6], focus on ex-post measurement of abnormal returns around predetermined event dates. While highly effective in evaluating market reactions, such approaches do not provide ex-ante probabilities regarding the recurrence of events. Bayesian change-point detection methods, along with Markov regime-switching models, are excellent for identifying structural breaks or hidden patterns in data after the fact. Typically, the entire dataset is analyzed to make these assessments. Although they are very effective for dividing data into segments, they are not always the best choice for predicting when specific events will occur in real time.

Volatility models, such as the ARCH framework, effectively capture volatility clustering and leverage effects, thereby providing accurate forecasts of variance. Although these models effectively demonstrate dispersion patterns, they are not appropriate for forecasting the precise likelihood of extreme events occurring at specific times.

On the other hand, extreme value theory helps us understand both the frequency and severity of rare events; however, it does not address when these events are likely to occur. Duration and hazard models provide a probabilistic structure for event occurrence but typically rely on static parametric hazard functions, mainly exponential or Weibull. They are mainly used for firm defaults or credit events over fixed time horizons. Likewise, jump-diffusion and stochastic volatility models with jumps account for discontinuities in asset prices but treat jump intensities as model parameters rather than directly estimating event recurrence probabilities from data. Self-exciting processes, following Hawkes, introduce temporal clustering through excitation kernels but generally assume a fixed or slowly varying baseline.

While prior studies have relied on Hawkes processes within unified parametric structures to model cryptocurrency return dynamics, the present study provides new evidence that separating event identification from probabilistic modeling leads to a qualitatively different view of extreme market behavior. By identifying extreme events independently and introducing Hawkes-type excitation only as a short-run probability adjustment, the analysis indicates that clustering effects manifest predominantly in the timing of event recurrence, rather than in the magnitude of returns. This finding offers new evidence that self-excitation in cryptocurrency markets is more appropriately interpreted as a mechanism governing temporal dependence in event probabilities rather than as a generator of extreme price movements per se. From this perspective, the role of Hawkes processes shifts away from structural explanation toward calibrated, forward-looking probability forecasting, thereby complementing existing Hawkes-based approaches and enhancing their relevance for real-time risk assessments.

5.4. Limitations and Directions for Future Research

First, using the methodology on Bitcoin returns does not imply that the distribution of the events for all other time series is identical. Therefore, rigorous testing of this methodology is necessary. Second, the definition of events or the methodology for selecting events could be modified to yield different results while incorporating exogenous variables such as trading volume and current market sentiment. This approach can utilize both automatic and manual event selection to achieve better results. Third, in this study, we demonstrated the methodology by utilizing a set of fixed parameters. Adjusting the methodology's parameters, such as time intervals, outlier contamination factors, weights, and half-life parameters, could provide different probabilities. The weight and decay parameters in this study are constant; however, the volatile nature of financial markets could require the use of dynamic parameters that adjust to the current time series structures.

Additionally, the low event count makes it difficult to calibrate an accurate probability function; more data would likely lead to more reliable results. Another limitation is that the methodology does not address the issue of predicting the scale of an incoming event and calculating the scale probabilities, which could introduce risk-adjusted probability. Extensions to the model could include seasonality and threshold components, a separation of the excitation factor into fast and slow components, and even a multivariate system in which event risk spillovers could influence the probabilities. Finally, the methodology could be refined to model the magnitude and probabilities of extreme market returns more accurately, potentially through the incorporation of tail-weighted distributions.

5.5. Implications

This framework could be practically used by risk managers, hedge funds, and traders to assess the risks of both positive and negative events and adapt their strategies accordingly.

By having the estimated probabilities of unusual returns in the asset prices, they can either open or close trades, readjust leverages, and reprogram algorithmic trading rules.

Regarding academic significance, this framework expands event mapping from an empirical distribution to a dynamic parametric estimation, rather than focusing on static event distribution studies. The methodology calculates the real-time probability and refines it with an excitation factor, differentiating it from other methods in the literature.

6. Conclusions

This study introduced a probabilistic framework for detecting and forecasting extreme return events in financial time series using daily Bitcoin data as an empirical setting. This approach combines machine learning event identification, empirical modeling of inter-arrival times, Hawkes-style excitation dynamics, and logistic regression to generate real-time probability estimates. Through these components, the framework provides a structured and transparent method for quantifying the likelihood of future positive and negative extreme events.

The methodology was explicitly built around four integrated steps: (i) machine-learning anomaly detection to identify extreme return days; (ii) estimation of empirical inter-arrival distributions to establish a dynamic baseline probability; (iii) incorporation of short-run clustering effects through a 7-day Hawkes excitation kernel; and (iv) a logistic regression layer that merges long-run and short-run information into real-time forecasts. This explicit sequence clarifies the mechanics of the model and removes the need for the reader to infer the methodological details.

The empirical findings show that positive extreme events occur more frequently than negative ones and that forecast accuracy consistently improves as more event data accumulate. Brier scores and log loss values both exhibited downward trends, indicating progressive calibration and improved discrimination over time. These results highlight the adaptability of the framework and its capacity to incorporate an expanding stream of observations.

Compared with previous methodologies, such as volatility-based models (e.g., GARCH), tail-distribution approaches (Extreme Value Theory), or classical Hawkes processes, this framework differs in three key ways: (i) it is not based on exogenous thresholds or variance-only proxies, but instead identifies events endogenously through anomaly detection; (ii) it does not assume a static baseline intensity, as the inter-arrival distribution updates dynamically; and (iii) it integrates detection, timing, and excitation into a unified structure rather than treating them as isolated modeling tasks. This distinction clarifies how the proposed model extends the existing literature and what conceptual gap it fills.

The significance of this framework extends beyond academic interest. Investors, quantitative analysts, and risk managers can use such real-time probability forecasts to anticipate periods of heightened volatility, adjust their portfolio exposure, or refine their hedging strategies. Policymakers and market supervisors may also rely on early warning signals that indicate the build-up of instability to assist in systemic risk monitoring and intervention planning. By linking event detection and event recurrence in a single model, this framework provides practical value for both private and institutional stakeholders.

Overall, this study demonstrates that combining anomaly detection, empirical timing patterns, and excitation dynamics can yield a flexible and operational tool for forecasting extreme events in financial markets. Future research may extend this framework to multivariate settings, alternative asset classes, or intraday horizons and explore deeper integration with neural point-process models.

Author Contributions: Conceptualization, K.P. and I.K.; methodology, K.P. and I.K.; software, K.P.; validation, K.P.; formal analysis, K.P.; investigation, K.P. and I.K.; resources, K.P.; data curation, K.P.;

writing—original draft preparation, K.P. and I.K.; writing—review and editing, K.P., I.K. and O.P.; visualization, K.P.; supervision, O.P.; project administration, O.P.; All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Price data were acquired from Coingecko.com. The results can be found in: <https://doi.org/10.5281/zenodo.17973524>.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ARCH Autoregressive Conditional Heteroscedasticity
DDF Dynamic Distribution Fitting

Appendix A

This section contains a summary of the sensitivity analysis of the methodology using different Isolation Forest parameters.

In Table A1, we observe mean Brier scores ranging from 0.21 to 0.29 and mean log loss values from 0.69 to 1.10. Higher contamination factors led to more outliers, lower Brier scores, and lower mean log loss values.

Table A1. Positive Events Sensitivity Analysis: Performance Metrics.

Model Configuration [Estimators, Contamination]	Positive Events	Mean Brier Score	Mean Log Loss
[2, 0.05]	178	0.29	0.86
[50, 0.01]	66	0.23	0.80
[50, 0.05]	139	0.21	0.69
[100, 0.01]	68	0.26	0.97
[100, 0.05]	141	0.22	0.71
[200, 0.01]	63	0.24	1.10
[200, 0.05]	135	0.22	0.71

Table A2 shows mean Brier scores of 0.23 to 0.33 and log loss values of 0.99 to 1.17. Raising the contamination rate increases the number of flagged outliers, but unlike in the positive-event sensitivity analysis, this also drives both scores upward.

Table A2. Negative Events Sensitivity Analysis: Performance Metrics.

Model Configuration [Estimators, Contamination]	Negative Events	Mean Brier Score	Mean Log Loss
[2, 0.05]	84	0.23	1.17
[50, 0.01]	99	0.28	1.02
[50, 0.05]	163	0.30	1.08
[100, 0.01]	98	0.28	1.01
[100, 0.05]	159	0.31	1.13
[200, 0.01]	110	0.27	0.99
[200, 0.05]	167	0.33	1.17

Table A3 shows that varying the Isolation Forest parameters results in insignificant correlations between the absolute tail returns and the combined probability values.

Table A3. Extreme Returns Correlation Sensitivity Analysis.

Model Configuration [Estimators, Contamination]	Pearson-P.T.	Spearman-P.T.	Pearson-N.T.	Spearman-N.T.
[2, 0.05]	−0.12	−0.09	−0.15	−0.10
[50, 0.01]	−0.07	−0.10	−0.14	−0.09
[50, 0.05]	−0.13	−0.15 *	−0.08	0.00
[100, 0.01]	−0.08	−0.12	−0.14	−0.10
[100, 0.05]	−0.11	−0.11	−0.13	−0.03
[200, 0.01]	−0.07	−0.10	−0.12	−0.07
[200, 0.05]	−0.12	−0.12	−0.12	−0.01

Note. For the Positive Tail (P.T.) results, the number of observations was 120, and for the Negative Tail (N.T.) results, the number of observations was 107. The asterisk (“*”) denotes significance at the 10% level.

References

- Corrado, C.J. Event Studies: A Methodology Review. *Account. Financ.* **2011**, *51*, 207–234. [\[CrossRef\]](#)
- Anderson, D.N.; Randall, G.E.; Whitaker, R.W.; Arrowsmith, S.J.; Arrowsmith, M.D.; Fagan, D.K.; Taylor, S.R.; Selby, N.D.; Schult, F.R.; Kraft, G.D.; et al. Seismic Event Identification. *WIREs Comput. Stat.* **2010**, *2*, 414–432. [\[CrossRef\]](#)
- Thompson, S.; Kaptoge, S.; White, I.; Wood, A.; Perry, P.; Danesh, J. The Emerging Risk Factors Collaboration Statistical Methods for the Time-to-Event Analysis of Individual Participant Data from Multiple Epidemiological Studies. *Int. J. Epidemiol.* **2010**, *39*, 1345–1359. [\[CrossRef\]](#)
- Martinez, J.C. Methodology for Conducting Discrete-Event Simulation Studies in Construction Engineering and Management. *J. Constr. Eng. Manag.* **2010**, *136*, 3–16. [\[CrossRef\]](#)
- Ogasawara, E.; Salles, R.; Porto, F.; Pacitti, E. *Event Detection in Time Series*; Synthesis Lectures on Data Management; Springer Nature: Cham, Switzerland, 2025.
- MacKinlay, A.C. Event Studies in Economics and Finance. *J. Econ. Lit.* **1997**, *35*, 13–39.
- Barry, D.; Hartigan, J.A. A Bayesian Analysis for Change Point Problems on JSTOR. Available online: <https://www.jstor.org/stable/2290726?seq=1> (accessed on 6 October 2025).
- Fearnhead, P. Exact and Efficient Bayesian Inference for Multiple Changepoint Problems. *Stat. Comput.* **2006**, *16*, 203–213. [\[CrossRef\]](#)
- Hamilton, J.D. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* **1989**, *57*, 357–384. [\[CrossRef\]](#)
- Engle, R. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. Available online: <https://www.jstor.org/stable/1912773?seq=1> (accessed on 6 October 2025).
- Bollerslev, T. Generalized Autoregressive Conditional Heteroskedasticity. *J. Econom.* **1986**, *31*, 307–327. [\[CrossRef\]](#)
- Nelson, D.B. Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica* **1991**, *59*, 347–370. [\[CrossRef\]](#)
- Zakoian, J.-M. Threshold Heteroskedastic Models. *J. Econ. Dyn. Control* **1994**, *18*, 931–955. [\[CrossRef\]](#)
- Ait-Sahalia, Y.; Fan, J.; Li, Y. The Leverage Effect Puzzle: Disentangling Sources of Bias at High Frequency. *J. Financ. Econ.* **2013**, *109*, 224–249. [\[CrossRef\]](#)
- Embrechts, P.; Klüppelberg, C.; Mikosch, T. *Modelling Extremal Events*; Springer: Berlin/Heidelberg, Germany, 1997.
- Cox, D.R. Regression Models and Life-Tables. *J. R. Stat. Society. Ser. B (Methodol.)* **1972**, *34*, 187–220. [\[CrossRef\]](#)
- Shumway, T. Forecasting Bankruptcy More Accurately: A Simple Hazard Model. Available online: <https://www.jstor.org/stable/10.1086/209665?seq=1> (accessed on 6 October 2025).
- Merton, R.C. Option Pricing When Underlying Stock Returns Are Discontinuous. *J. Financ. Econ.* **1976**, *3*, 125–144. [\[CrossRef\]](#)
- Bates, D. Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. Available online: <https://www.jstor.org/stable/2962366?seq=1> (accessed on 7 October 2025).
- Hawkes, A.G. Spectra of Some Self-Exciting and Mutually Exciting Point Processes. *Biometrika* **1971**, *58*, 83–90. [\[CrossRef\]](#)
- Phillips, R.C.; Gorse, D. Mutual-Excitation of Cryptocurrency Market Returns and Social Media Topics. In *Proceedings of the 4th International Conference on Frontiers of Educational Technologies*; Association for Computing Machinery: New York, NY, USA, 2018; pp. 80–86.
- Luo, R.; Krishnamurthy, V.; Blasch, E. Hawkes Process Modeling of Block Arrivals in Bitcoin Blockchain. *arXiv* **2022**, arXiv:2203.16666.
- Cestari, R.G.; Barchi, F.; Busetto, R.; Marazzina, D.; Formentin, S. Hawkes-Based Cryptocurrency Forecasting via Limit Order Book Data. *arXiv* **2023**, arXiv:2312.16190.
- Yang, S.Y.; Pirjol, D.; Zhang, B.; Li, Q. Cryptocurrency Jump Contagion with Market Sentiment Events: A Study of High Frequency Cross Effect. *Eur. J. Financ.* **2025**, 1–19. [\[CrossRef\]](#)

25. Kaminsky, G.L.; Reinhart, C.M. The Twin Crises: The Causes of Banking and Balance-Of-Payments Problems. *Am. Econ. Rev.* **1999**, *89*, 473–500. [[CrossRef](#)]
26. Borio, C.E.V.; Drehmann, M. Assessing the Risk of Banking Crises—Revisited. *BIS Q. Rev.* **2009**, 29–46. Available online: <https://ssrn.com/abstract=1513316> (accessed on 29 January 2026).
27. Liu, F.T.; Ting, K.M.; Zhou, Z.-H. Isolation Forest. In Proceedings of the 2008 Eighth IEEE International Conference on Data Mining, Pisa, Italy, 15–19 September 2008; pp. 413–422.
28. Malhotra, P.; Vig, L.; Shroff, G.; Agarwal, P. Long Short Term Memory Networks for Anomaly Detection in Time Series. In Proceedings of the 23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN), Bruges, Belgium, 22–24 April 2015; Volume 89, p. 94.
29. Foorthuis, R. On the Nature and Types of Anomalies: A Review of Deviations in Data. *Int. J. Data Sci. Anal.* **2021**, *12*, 297–331. [[CrossRef](#)]
30. Zhu, L.; Gao, R.; Huang, C.; Shi, Q.; Shi, Z. A Dynamic Event-Based Recursive State Estimation for Mobile Robot Localization. *Electronics* **2024**, *13*, 3227. [[CrossRef](#)]
31. Gao, Y.; Li, Y.; Hua, Z.; Chen, J.; Wu, Y. A Dynamic Event-Triggered Secure Monitoring and Control for a Class of Discrete-Time Markovian Jump Systems: A Plug-and-Play Architecture. *Information* **2024**, *15*, 649. [[CrossRef](#)]
32. Phuntsho, T.; Gonsalves, T. Maximizing Net Present Value for Resource Constraint Project Scheduling Problems with Payments at Event Occurrences Using Approximate Dynamic Programming. *Algorithms* **2024**, *17*, 180. [[CrossRef](#)]
33. Gecko Labs Pte. Ltd. CoinGecko: Cryptocurrency Prices, Charts, and Crypto Market Cap. Available online: <https://www.coingecko.com/> (accessed on 4 September 2025).
34. Python Software Foundation. *Python*, 3.10.9; Python Software Foundation: Wilmington, DE, USA, 2022. Available online: <https://www.python.org> (accessed on 29 January 2026).
35. Harris, C.R.; Millman, K.J.; van der Walt, S.J.; Gommers, R.; Virtanen, P.; Cournapeau, D.; Wieser, E.; Taylor, J.; Berg, S.; Smith, N.J.; et al. Array Programming with NumPy. *Nature* **2020**, *585*, 357–362. [[CrossRef](#)] [[PubMed](#)]
36. The Pandas Development Team. *Pandas-Dev/Pandas: Pandas*; Zenodo: Geneva, Switzerland, 2020. [[CrossRef](#)]
37. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-Learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
38. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [[CrossRef](#)] [[PubMed](#)]
39. Cokelaer, T.; Kravchenko, A.; Lahdjirayhan; Msat59; Ferrari, V.; Varma, A.; L, B.; Stringari, C.E.; Brueffer, C.; Broda, E.; et al. *Cokelaer/Fitter*, V1.7.1 2024; Zenodo: Geneva, Switzerland, 2024. [[CrossRef](#)]
40. Vrieze, S.I. Model Selection and Psychological Theory: A Discussion of the Differences between the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). *Psychol. Methods* **2012**, *17*, 228–243. [[CrossRef](#)] [[PubMed](#)]
41. Snipes, M.; Taylor, D.C. Model Selection and Akaike Information Criteria: An Example from Wine Ratings and Prices. *Wine Econ. Policy* **2014**, *3*, 3–9. [[CrossRef](#)]
42. Brier, G.W. Verification of Forecasts Expressed in Terms of Probability. *Mon. Weather. Rev.* **1950**, *78*, 1–3. [[CrossRef](#)]
43. Bishop, C.M.; Nasrabadi, N.M. *Pattern Recognition and Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2006; Volume 4.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.