Journal of Information and Data Management, 2026, 17:1, doi: 10.5753/jidm.2026.5728

HDP+: Leveraging Anomaly and Change Point Detection for Pump-and-Dump in Cryptocurrency Exchanges

Matheus S. Moura [**Federal Center for Technological Education of Rio de Janeiro - CEFET/RJ** | ***matheus.moura@aluno.cefet-rj.br***]
Laís Baroni [**Federal Center for Technological Education of Rio de Janeiro - CEFET/RJ** | ***lais.baroni@aluno.cefet-rj.br***]
Eduardo Ogasawara [**Federal Center for Technological Education of Rio de Janeiro - CEFET/RJ** | ***eogasawara@ieee.org***]
Diogo S. Mendonça [**Federal Center for Technological Education of Rio de Janeiro - CEFET/RJ** | ***diogo.mendonca@cefet-rj.br***]

Federal Center for Technological Education of Rio de Janeiro - CEFET/RJ, Av. Maracanã, 229 - Maracanã – Rio de Janeiro/RJ - CEP: 20271-110, Brazil.

Received: 23 March 2025 • **Published:** 13 March 2026

Abstract Cryptocurrencies are increasingly gaining relevance in financial markets, attracting both retail and institutional investors. As a result, assessing the risks associated with these new assets has become more important. Unlike traditional market assets, which are regulated and centrally managed, cryptocurrencies were designed with an egalitarian nature, enabling peer-to-peer transactions and providing varying levels of anonymity. These characteristics contribute to a favorable environment for illicit activities, including market manipulation. One such manipulation technique is the pump-and-dump (PD) scheme, which exploits the decentralized and anonymous nature of cryptocurrency markets. Typically orchestrated through public groups on social media platforms, these schemes involve coordinated surges in buy orders to artificially inflate a coin's price, followed by rapid sell-offs that leave unsuspecting investors with losses. The detection of PD schemes is important because they compromise the integrity and reputation of cryptocurrency markets, undermining investor confidence and contributing to market instability. Most prior research on PD detection has focused on anomaly detection techniques. In this study, we investigate whether combining anomaly detection with change point detection in a hybrid framework can enhance detection performance. We propose HDP+, an improved version of the original HDP method, which processes raw trading records from exchanges and applies an ensemble of anomaly detection and change point detection algorithms. The primary distinction between HDP and HDP+ lies in the anomaly detection component: while HDP relies on traditional volatility-based methods, HDP+ analyzes the time series of rush orders. Experiments conducted on a dataset of confirmed PD events yielded results of 96.4% precision, 89.3% recall, and a 92.7% F1-score, surpassing previous statistical approaches to PD detection.

Keywords: Cryptocurrency, Pump-and-dump, Fraud detection, Anomaly detection, Change point detection

1 Introduction

The use of cryptocurrencies has increased over the past decade, resulting in market expansion across traditional financial sectors and in public perception [Jalal *et al.*, 2021]. This expansion is illustrated by events such as the approval of the first Bitcoin futures Exchange-Traded Fund (ETF) by the U.S. Securities and Exchange Commission (SEC) in 2021 [Wursthorn, 2021], and the initiation of spot ETF trading on American stock exchanges in early 2024 [Schmitt, 2024]. Public interest has followed periods of rapid market growth, such as those observed in late 2017 and early 2018, which were widely covered by mainstream media and tend to recur during subsequent market rallies [Steinmetz *et al.*, 2021]. As a result, studies focused on assessing the risks associated with this asset class have become necessary.

Most cryptocurrencies operate without centralized control and support a certain degree of anonymity. These characteristics facilitate their use in illicit practices, including money laundering, drug trafficking, and cyberattacks [Kethineni and Cao, 2020]. One notable type of fraudulent practice involving cryptocurrencies is the pump-and-dump (PD) scheme, in which an asset's price is artificially increased and then the asset is sold at a higher price [Victor and Hagemann, 2019].

PD schemes in cryptocurrency markets are often organized through public groups on platforms that offer encrypted communication and anonymity, such as Telegram or Discord. These groups recruit members until they reach sufficient financial capacity to influence asset prices. The scheme begins with a pre-announcement made a few days in advance, which includes information about the exchange[1], the pairing coin[2], and the exact timing of the pump [La Morgia *et al.*, 2020]. At the scheduled time, the targeted asset is disclosed, and members begin coordinated buying, holding, and promotion activities to increase the price. The pump usually reaches its peak within minutes, followed by rapid sell-offs at the first sign of decline. The asset's price and trading volume typically return to pre-pump levels within approximately thirty minutes [Xu and Livshits, 2019].

The rapid and artificial price increases induced by PD schemes deviate significantly from the typical patterns observed in cryptocurrency time series. Existing studies primarily rely on anomaly detection methods, yet these approaches often have problems with precision and adaptability [Rajaei and Mahmoud, 2023]. To address these short-

[1] An organized market or center for trading cryptocurrencies.
[2] Cryptocurrency used to trade against another cryptocurrency.

comings, this study explores the effectiveness of integrating anomaly detection with change point detection techniques, evaluating whether this hybrid approach improves the detection performance of PD events compared to approaches that rely solely on anomaly detection.

HD Pump Plus (HDP+) is a hybrid detection method designed to identify PD schemes in cryptocurrency markets by integrating anomaly detection (AD) and change point detection (CPD) techniques. The method processes raw trading records into time series representations that capture essential trading attributes such as price, volume, and rush orders. This work is an extension of the Moura *et al.* [2024] HD Pump (HDP), enhancing the method through the incorporation of the rush order feature, which identifies rapid bursts of trading activity within the same millisecond—a pattern associated with PD scheme organizers. The workflow includes data preprocessing steps, such as applying the cumulative sum (CUSUM) transformation to the volume series, allowing CPD methods to detect abrupt structural changes.

The detection process in HDP+ is structured into two main components: the AD component, which utilizes Robust Empirical Mode Decomposition (REMD) [Souza *et al.*, 2024] to identify anomalies in the rush order series; and the CPD component, which employs an ensemble of the Chow Test and GFT on the CUSUM volume series to detect structural changes [Ogasawara *et al.*, 2025]. The final classification of PD events is determined by the intersection of positive labels from both components, prioritizing the precision metric. By leveraging the strengths of both AD and CPD, HDP+ achieves a performance improvement without requiring model training, making it a robust solution even when compared to more complex methods.

The experimental evaluation of HDP+ was conducted using a dataset of 178 confirmed PD events in the Binance exchange[3], obtained from the dataset provided by La Morgia *et al.* [2020]. The results, summarized in Table 4, demonstrate a significant performance improvement of HDP+ over the original HD Pump and other methods. It achieves a precision of 96.4%, recall of 89.3%, and F1-score of 92.7%. Compared to the original HD Pump, HDP+ significantly improves precision and recall, increasing these metrics by 39.6% and 5%, respectively. These improvements enhance the detection and prevention of PD schemes by expanding the range of detection frameworks and addressing practices that compromise the integrity and reputation of cryptocurrency markets. The method balances high precision and recall, outperforming prior statistical approaches in the identification of PD events.

The remainder of this paper is organized as follows. Sections 2 and 3 provide the theoretical background and review related work. Section 4 details the HDP+ method. Section 5 outlines the experimental protocol and presents the results. Finally, Section 6 concludes the paper and discusses potential future research directions.

2 Background

The PD scheme has a long history in the stock market and a straightforward premise [Kamps and Kleinberg, 2018]. Initially, the perpetrators identify a publicly traded security as their target, typically small-cap stocks or low-liquidity assets due to their susceptibility to price manipulation, allowing them to exert significant influence with relatively small investments. Subsequently, they accumulate significant quantities of this security. Following the acquisition, they promote the security aggressively through misleading information and market hype, disseminating false and deceptive information aimed at artificially inflating its price [Kramer, 2005].

Unlike traditional financial systems, most cryptocurrencies operate without central authorities to process transactions, instead utilizing the peer-to-peer technology known as blockchain. The egalitarian nature of cryptocurrencies has introduced a new era of decentralized authority, providing transaction privacy, anonymity, and a lack of deterrence [Kethineni and Cao, 2020]. These characteristics have created a favorable environment for online criminal activities using cryptocurrencies, such as PD schemes.

A PD scheme in the cryptocurrency market is a form of market manipulation in which the price of a cryptocurrency is artificially inflated to generate profits for the scheme insiders. This practice relies on coordinated efforts to create misleading hype, luring unsuspecting investors into purchasing the targeted asset at an inflated price. Once the price peaks, the perpetrators sell off their holdings, causing a sharp decline that leaves late investors with significant losses [Kamps and Kleinberg, 2018].

The scheme follows a structured process, illustrated in Figure 1. It begins with the group bootstrap, where organizers establish groups on platforms like Telegram or Discord to gather participants. In some cases, exclusive VIP groups receive privileged early access to pump information in exchange for a fee. Following this, an announcement is made, providing details of the planned pump, including the date, time, exchange, and the cryptocurrency to be targeted [Hu *et al.*, 2023]. Before the event, the group organizers accumulate the asset, discreetly purchasing large quantities of the chosen cryptocurrency to maximize their eventual gains [Xu and Livshits, 2019].

At the designated time, the pump phase begins, during which organizers publicly disclose the targeted asset and urge members to buy. This coordinated activity creates a surge in demand, driving the price upward. Misleading information may be disseminated to attract external investors [La Morgia *et al.*, 2020] to amplify the effect. Once the price peaks, the dump phase ensues, where the group members sell off their holdings, triggering a rapid price collapse. Finally, in the post-pump phase, organizers often share misleading profitability statistics to maintain the illusion of success and recruit participants for future schemes [Xu and Livshits, 2019].

The detection of PD schemes in cryptocurrency markets has been studied primarily with classification-based anomaly detection methods, including statistical and machine learning-based techniques [Rajaei and Mahmoud, 2023]. Both approaches focus on identifying unusual trading patterns within the time series data of the targeted cryp-

[3] The largest cryptocurrency exchange at the time of this work.

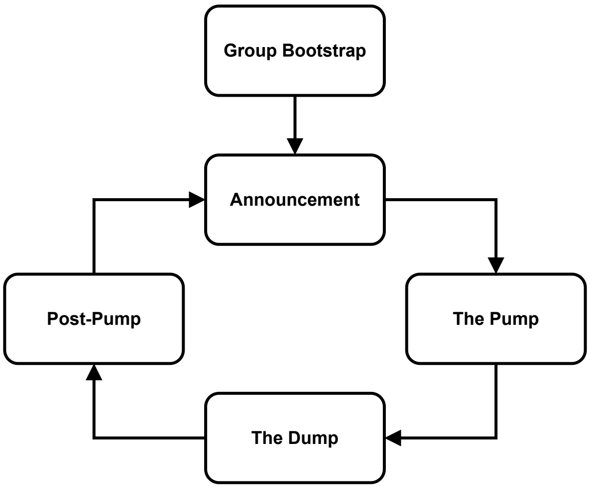

Figure 1. Illustration of the stages of a pump-and-dump (PD) scheme in cryptocurrency markets. The process begins with recruiting and organizing a PD group and then announcing the pump details. During the pump phase, coordinated buying artificially inflates the price, often attracting unsuspecting investors. Once the price peaks, the dump starts, where organizers sell their holdings, causing a price collapse. Finally, in the post-pump phase, misleading statistics may be shared to maintain the illusion of profitability and recruit new participants.

tocurrency. The labeling of PD events is typically conducted through longitudinal studies that monitor PD groups on social media platforms, such as the one conducted by La Morgia *et al.* [2020]. Given the structured nature of PD schemes, their detection requires methods capable of identifying anomalous price and volume patterns within time series data.

Time series analysis is a widely used tool in financial markets, enabling researchers to examine price and volume dynamics over time [Han *et al.*, 2022]. Since PD schemes introduce abrupt fluctuations in asset prices and trading volumes, time series analysis is important for identifying these irregularities. A time series is defined as a sequence of ordered observations recorded at specific intervals, represented as $(z_1, z_2, \ldots, z_t)$, where z_t $(t = 1, \ldots, T)$ denotes the observation at a given time step, and T represents the total length of the series. Examples include hourly cryptocurrency price fluctuations and minute-by-minute trade volumes. The analysis of time series data facilitates the identification of trends, patterns, and structural changes over time [Ogasawara *et al.*, 2025], as well as the detection of anomalies such as PD schemes.

PD schemes create abrupt and significant price and trading volume fluctuations that do not conform to historical trends [Kamps and Kleinberg, 2018]. In order to detect the PD events, most studies employ anomaly detection methods [Rajaei and Mahmoud, 2023]. Anomalies in time series are abnormal patterns or events, that appear not to be generated by the same process that generated most of the observations of the series. Formally, an observation z_t is considered an anomaly if its temporal component deviates significantly from both its expected value based on the k preceding and k following observations. This condition is satisfied when both deviations exceed a predefined threshold σ [Ogasawara *et al.*, 2025].

For example, a basic method for anomaly detection in time series is the application of statistical thresholds to moving averages of price and trading volume. Data points exceeding three standard deviations from the series mean can be flagged as potential anomalies [Ogasawara *et al.*, 2025]. Despite the effectiveness of statistical threshold-based anomaly detection, distinguishing PD events from normal market volatility remains challenging. More advanced techniques, such as machine learning-based methods or hybrid approaches combining anomaly detection with change point detection, have been explored to address these limitations.

3 Related Work

Several approaches have been proposed for detecting PD schemes in cryptocurrency markets, ranging from statistical anomaly detection to machine learning and deep learning models. Kamps and Kleinberg [2018] was one of the pioneering studies on cryptocurrency PD schemes. It offered an early formalization of these schemes and presented a methodology for detecting them using AD algorithms. Their method identifies local conditional point anomalies by detecting simultaneous price and volume trend irregularities. Furthermore, their method has parameters that optimize recall, precision, or a balance between metrics. However, a limitation of their study lies in the dataset used: it lacks confirmation of whether the analyzed events were actual PD schemes.

The limitation of labeled data encountered by Kamps and Kleinberg [2018] was subsequently addressed by the work of La Morgia *et al.* [2020], who introduced a publicly available dataset of confirmed PD schemes. Their work focused on monitoring PD groups that primarily operated on the Binance exchange, with order data restored using the Binance API, which allows retrieval of every transaction within the complete trading pair history. La Morgia *et al.* [2020] also reproduced the work of Kamps and Kleinberg [2018] using their dataset, and the reported results can be found in Table 4.

Building upon hybrid detection methods, HDP+ extends the HD Pump method introduced by Moura *et al.* [2024], incorporating rush order-based anomaly detection to enhance PD detection accuracy. Both methods share the same overall workflow structure and CPD component. However, they differ in the AD component. While HDP+ incorporates rush order-based anomaly detection, the original HDP detects volatility anomalies in price time series. Specifically, the AD component in HDP employs an ensemble detector combining GARCH and RED techniques from Salles *et al.* [2020], applied to the first-order differentiated price series to identify fluctuations indicative of PD events.

Beyond the studies compared in this work, there are more complex methods in the literature that focus on the detection of PD schemes in cryptocurrency markets, particularly through machine learning-based approaches [Rajaei and Mahmoud, 2023]. An example is the work by La Morgia *et al.* [2023], which expanded upon La Morgia *et al.* [2020] by introducing another type of PD scheme: the crowd pump. Unlike traditional PD schemes, crowd pumps involve large-scale, decentralized coordination, often through social media platforms. A notable example is the GameStop short squeeze of 2021, which La Morgia *et al.* [2023] used as a case study to extend their previous work. Their detection of the fraud schemes achieved a detection performance of 98.2% precision and 91.2% recall, using a random forest model that leverages a rush orders feature. Section 4 presents a detailed discussion of rush orders. However, despite their effectiveness, rush orders have limitations, as their utility relies on access to complete trading records and characteristics of the data from the Binance exchange, which may not be available for every exchange.

While previous studies primarily relied on statistical and machine learning-based anomaly detection, Bello *et al.* [2023] introduced a deep learning approach. Their method leverages an LSTM-based autoencoder trained on Bitcoin valuations to predict the valuations of cryptocurrencies. Detection is triggered through thresholding based on a Gaussian tail condition. Unlike prior studies that focus on post hoc detection, Bello *et al.* [2023] emphasize real-time detection, enabling most PD events to be identified within five minutes using one-minute resolution data. The model achieved a precision of 78%, a recall of 83%, and an F1-score of 80%.

We adopt the work of Kamps and Kleinberg [2018] as a benchmark primarily because it employs the same time resolution as our method and is evaluated using the same dataset provided by La Morgia *et al.* [2020], ensuring a fair comparison. Although La Morgia *et al.* [2020] and La Morgia *et al.* [2023] both utilize rush order-based features, differences in time resolution limit direct comparison with our approach. In contrast, the method proposed by Bello *et al.* [2023] is designed for low-latency detection of PD schemes and is evaluated on a distinct dataset, differentiating it from our approach.

4 HD Pump Plus (HDP+)

HDP+ builds upon the original HD Pump method proposed by Moura *et al.* [2024], which represented an initial effort to construct a hybrid detector by integrating anomaly detection and change point detection techniques for identifying PD events in cryptocurrency exchanges. Both HDP and HDP+ follow a similar workflow, as illustrated in Figure 2. In this section, we first provide an explanation of the HDP workflow. Subsequently, we describe the enhancements introduced in HDP+, with a focus on the features that distinguish it from the HDP.

The workflow begins by transforming input data from cryptocurrency exchanges into time series, enabling the application of event detection methods. Both methods focus on data from the Binance exchange[4], which, at the time of this study, was the largest cryptocurrency exchange. This choice is motivated by the expectation that PD schemes on Binance require the coordination of larger trading volumes. Moreover, Binance also provides comprehensive access to historical trade records through its API, enabling detailed analyses of past PD events.

The input data obtained from the exchanges is in the format of trading records. Which are the chronological sequence of executed orders on the exchange platform within a given time interval. As shown in Table 1, each record typically includes the trading pair—comprising the target cryptocurrency and its pairing asset, which is not necessarily a cryptocurrency but is most commonly Bitcoin (BTC)—along with the timestamp of execution, the traded volume denominated in the pairing asset, the quantity of the cryptocurrency exchanged, and the trade direction (buy or sell). The price of the traded asset is not always explicitly provided, as it can be derived by dividing the traded volume by the quantity.

However, the trading records must be transformed into time series to enable analysis. Therefore, the first step in the workflow is the transformation of these records into time series. This process involves aggregating records into fixed time intervals, referred to as chunks. In the HDP method, for each chunk, the average price and total traded volume are computed, resulting in two time series for the asset.

Besides that, the time series extraction step also enriches the data by converting price and volume from Bitcoin into US dollars. This is done to help mitigating the noise generated by cryptocurrencies high volatility market. A sample of the extracted time series is shown in Table 2.

Following the extraction of time series, the subsequent step in the workflow is preprocessing. This stage aims to enhance the performance and reliability of the detection algorithms by emphasizing patterns and structural changes within the data. The HDP method applies two distinct preprocessing techniques, each specifically designed to its corresponding detection component.

The preprocessing technique employed for the HDP AD component is first-order differencing (DIFF). This technique is applied to the average price time series to highlight fluctuations in volatility, thereby facilitating the identification of PD events as anomalies. Formally, given a time series z_t at time t, the DIFF series d_t is defined as $d_t = z_t - z_{t-1}$ [Shumway and Stoffer, 2017]. An example of the resulting time series after applying DIFF preprocessing is shown in Figure 3.

For the CPD component, HDP employs the CUSUM preprocessing technique. While PD events are typically analyzed using anomaly detection methods [Rajaei and Mahmoud, 2023], the application of CUSUM transforms the short-term effects of these schemes into structural changes in the time series, thereby enabling the use of change point detection algorithms.

As depicted in Figure 4, the CUSUM preprocessing converts each observation into the cumulative sum of itself and all preceding observations. Formally, if z_t represents the original time series at time t, the CUSUM series c_t is defined as $c_t = \sum_{i=1}^{t} z_i$. The CUSUM series emphasizes shifts in the level or mean of the series by cumulatively summing the data points, thereby aiding in the identification of significant alterations in the statistical properties of the series [Takeuchi and Yamanishi, 2006].

Following the preprocessing stage, the resulting time series are processed by the detection components of the HDP method. The AD component utilizes the DIFF average price series, while the CPD component operates on the CUSUM of the volume series. Both components employ off-the-shelf event detection algorithms provided by the Harbinger frame-

[4] `https://www.binance.com`

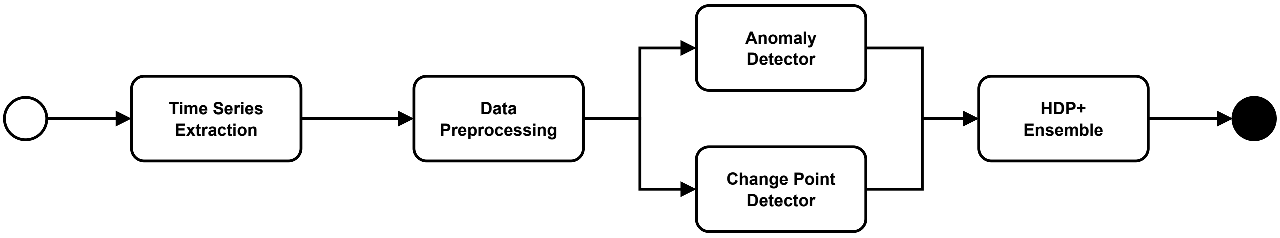

Figure 2. Workflow shared by both the HD Pump (HDP) and HD Pump Plus (HDP+) methods for the detection of pump-and-dump schemes in cryptocurrency exchanges.

Table 1. Sample of trading records for the Everex (EVX) cryptocurrency token paired with Bitcoin (BTC) on the Binance exchange.

Coin Pair	Datetime	Side	Price	Amount	Volume
EVX/BTC	2020-12-28T21:33:52.620Z	BUY	1.073e-05	3680.0	0.0394864
EVX/BTC	2020-12-28T21:33:56.493Z	SELL	1.068e-05	52.0	0.00055536
EVX/BTC	2020-12-28T21:33:58.183Z	SELL	1.068e-05	713.0	0.00761484

Table 2. Sample of the time series extracted for the Everex (EVX) cryptocurrency token, generated during the time series extraction step of the HD Pump method.

Bin	Number of Orders	Average Price (USD)	Volume Sum (USD)
2020-12-28 22:00:00	39	0.287319184922967	10791.8298698161
2020-12-28 23:00:00	40	0.291546664119495	9732.60237260589
2020-12-29 00:00:00	15	0.279370496295045	3442.74646652487

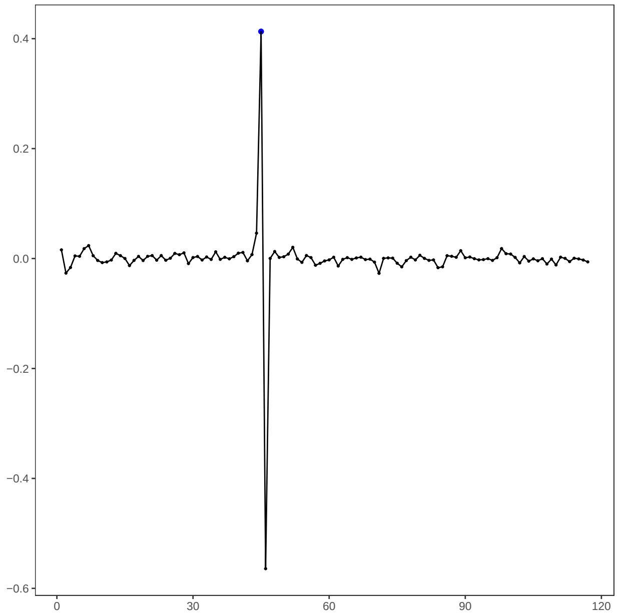

Figure 3. Visualization of the anomaly detection (AD) component of the HD Pump (HDP) method applied to the first-order differenced (DIFF) price time series of the pump-and-dump event involving the Everex cryptocurrency on September 1, 2021. The experiment was conducted using a chunk size of 3600 seconds and a total time frame of 180000 seconds.

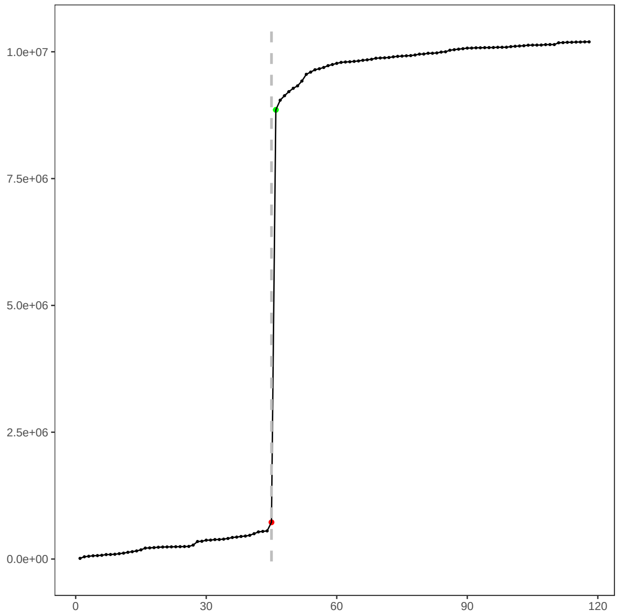

Figure 4. Visualization of the change point detection (CPD) component of the HD Pump (HDP) method applied to the cumulative sum (CUSUM) volume time series of the pump-and-dump event involving the Everex cryptocurrency on September 1, 2021. The experiment was conducted using a chunk size of 3600 seconds and a total time frame of 180000 seconds.

work[5] [Salles *et al.*, 2020].

Anomaly detection identifies abnormal patterns that deviate significantly from expected behavior, often indicating critical events within a domain. PD schemes create abrupt and unusual price and trading volume spikes, which can be modeled as anomalies within a time series [Kamps and Kleinberg, 2018]. Various techniques, including statistical models, machine learning, and signal decomposition methods detect

[5] Harbinger framework CRAN package: `https://cran.r-project.org/web/packages/harbinger/index.html`

these anomalies by analyzing deviations from typical series behavior [Ogasawara *et al.*, 2025].

The HDP AD component is designed to operate with high precision, based on the premise that PD events may be easy to detect as anomalies. However, the challenge lies in distinguishing genuine PD events from the high volatility of cryptocurrency markets. To address this, AD component prioritizes precision being able to detect easy PD events with high confidence, while letting a broader detection of events to the CPD component.

It is implemented as an ensemble of the Generalized Autoregressive Conditional Heteroscedasticity (GARCH) and the Harbinger RED algorithms. The GARCH algorithm is designed for identifying volatility anomalies in financial data time series. It operates by modeling the expected volatility at a given time point based on both prior and subsequent volatility. With an observation being classified as anomalous when its actual volatility deviates significantly from the GARCH's predicted value [Ogasawara *et al.*, 2025].

In contrast, the Harbinger RED algorithm is designed as an anomaly detection method that wraps the Empirical Mode Decomposition (EMD). The EMD is a time-frequency decomposition algorithm suited for nonlinear and nonstationary time series. The EMD decomposes the time series into various signals at different frequencies and detects the anomaly on the residual component [Ogasawara *et al.*, 2025]. The final result of the AD component is obtained by labeling as positive only those observations identified as anomalies by both the GARCH and RED algorithms.

Change point detection is the process of identifying specific observations where significant modifications occur in the statistical properties of a time series, such as shifts in mean, variance, or trend. These change points signal transitions between different underlying processes generating the time series and are crucial for system monitoring, adaptive modeling, and real-time decision-making [Ogasawara *et al.*, 2025]. The HDP method leverages CPD to identify PD schemes where abrupt volume fluctuations might indicate PD activity.

The HDP CPD component consists of an ensemble of the Chow test and the Generalized Fluctuation Test (GFT) algorithms. The Chow test is a regression-based change point detection method used to identify structural changes in a time series that can be modeled using linear regression [Ogasawara *et al.*, 2025]. It tests the null hypothesis of no structural change against the alternative hypothesis that regression coefficients vary over time. The test assumes a linear regression model, where observations are expressed as a function of independent variables with time-dependent coefficients. When the potential change point is known, the test involves fitting separate Ordinary Least Squares (OLS) regressions to the data segments before and after the change point, then computing an F-statistic to assess whether these segmented models provide a significantly better fit than a single regression model. Besides that, to improve the Chow test recall, the labeling process is adjusted such that any observation following a positively labeled instance is also assigned a positive label.

The GFT, in contrast, fits a regression model to the time series and evaluates the stability of its residuals under the null hypothesis of no structural change. Unlike the Chow test, which focuses on a single predefined change point, GFT detects structural changes by analyzing fluctuations in residuals or parameter estimates of a fitted regression model. Under the null hypothesis of no change, these fluctuations follow expected statistical boundaries derived from central limit theorems. Deviations beyond these limits indicate structural shifts in the time series.

The CPD ensemble strategy integrates both techniques by labeling an observation as a change point only when it is identified as such by both methods, thereby increasing the recall of the final classification. The final classification produced by the CPD component is determined through an ensemble approach, where an observation is labeled as positive if both the modified Chow test and the GFT method assign it a positive label.

The final stage of the workflow is the ensemble of the AD and CPD component results. Ensemble methods are widely used to improve detection performance by combining multiple base algorithms. The assumption is that individual detectors may perform better on specific subsets of inputs, and aggregating their outputs can enhance overall performance. Formally, given an input dataset X of n data tuples in a d-dimensional space, ensemble approaches build K base detectors, each operating on a randomly selected subspace $X_i \subseteq X$. Each base detector assigns an event score to the data tuples in its respective subspace. These scores are then aggregated to yield a final score for each tuple, which serves as the basis for the final result [Han *et al.*, 2022].

Finally, the output of the HDP method is the combination of the AD component and the CPD component results. Given the distinct characteristics of each component and the challenge of distinguishing true events from false positives caused by high market volatility, the ensemble strategy is designed to enhance precision without sacrificing recall. Specifically, if both components agree on at least one positively labeled observation, the final result consists of the intersection of their positive labels. Otherwise, the output defaults to the positive labels produced by the CPD component alone. This approach leverages on the high confidence detections of the AD component to filter false positives, without compromising the recall achieved by the CPD component.

On the other hand, HDP+ extends the original HDP method by incorporating the rush order feature introduced by La Morgia *et al.* [2020] and replacing the anomaly detection ensemble with the Robust Empirical Mode Decomposition (REMD) algorithm. Despite these enhancements, HDP+ retains the overall workflow illustrated in Figure 2, encompassing steps from time series extraction to the final result aggregation. However, since the detection components in HDP+ are no longer specialized for precision or recall, the final ensemble strategy is simplified accordingly.

Introduced by La Morgia *et al.* [2020], the rush order feature was a key component of their PD detection method, contributing significantly to its performance. This feature arises from limitations in the trading data provided by the Binance API, which does not include complete order-level information. The available Binance trading records are compressed, meaning that a single order executed at different prices is fragmented into multiple records. As a result, a single limit order[6] executed at multiple price levels is recorded as several individual trade records occurring at the same millisecond. The rush order feature is thus defined as the number of times trades are executed at the exact same millisecond.

Cryptocurrency exchanges try to emulate traditional brokerages by offering common order types, such as market and limit orders. Market orders execute immediately at the best available price, while limit orders execute only at a specified price or better. In PD schemes, organizers with prior knowledge can strategically place limit orders in advance to profit once the scheme begins. As illustrated in Figure 5, the rush order feature captures this behavior by identifying bursts of trades executed at the same millisecond, indicating the simultaneous triggering of pre-positioned limit orders.

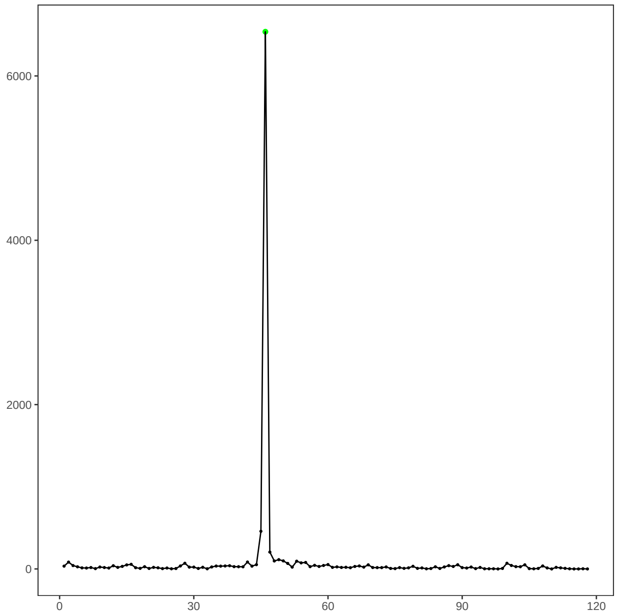

Figure 5. Visualization of the anomaly detection component of the HD Pump Plus (HDP+) method applied to the rush orders time series of the pump-and-dump event involving the Everex cryptocurrency on September 1, 2021. The experiment was conducted using a chunk size of 3600 seconds and a total time frame of 180000 seconds.

The rush order feature is highly effective, as highlighted in Table 3. Its contribution to performance improvement is further reinforced by the study conducted by La Morgia *et al.* [2023], discussed in Section 3, where it played a key role in achieving state-of-the-art results. The incorporation of rush orders also established a higher performance benchmark for the HD Pump workflow, as demonstrated by the results of HDP+ in Table 4. However, despite its advantages, the rush order feature presents limitations. Its applicability depends on the availability of complete trade records, which may not be accessible for all exchanges. Moreover, access to more detailed order type information could facilitate the development of an even more precise and discriminative feature, further enhancing the accuracy of PD detection.

Besides that, the HDP+ also introduces a new AD component in order to detect PD events in rush orders series. This new AD component relies on the Robust Empirical Mode Decomposition (REMD) algorithm, also from the Harbinger framework [Salles *et al.*, 2020]. REMD is a hybrid technique that integrates Empirical Mode Decomposition (EMD) with an Autoregressive Integrated Moving Average (ARIMA) model.

Table 3. Comparison of the rush order feature vs. first-order differentiation of the price series using a simple outlier detection method, which identifies events when an observation deviates by more than three standard deviations from the series mean.

Series	Precision	Recall	F1-score
Price diff	40.6%	69.1%	49.2%
Rush orders	72.0%	97.7%	80.2%

The EMD is used to decomposing the time series into a set of Intrinsic Mode Functions (IMFs) and employs a frequency separation technique based on roughness curve analysis to distinguish between high-frequency and low-frequency components [Souza *et al.*, 2024]. Subsequently, the ARIMA model is applied to the high-frequency components to generate time series predictions. Anomalies are then identified by computing the probability of an observation aligning with the ARIMA prediction. Observations classified as outliers within this probability distribution are flagged as anomalies [Souza *et al.*, 2024].

Finally, the HDP+ ensembles the results of this REMD-based AD component with the same CPD component of HDP. The AD component identifies anomalies within the rush order time series, which can be interpreted as abrupt increases in price triggering limit orders. These unusual spikes in rapid buying activity are characteristic of the initial pump phase of the PD scheme, as colluding members attempt to inflate the price quickly. Conversely, the CPD component is designed to identify significant structural changes in the cumulative volume series, detecting observations that exhibit notable shifts in overall trading volume. The ensemble of these results is determined by the intersection of the positive labels from both the AD and CPD components, ensuring that only the most confident positive labeled observations are retained.

5 Results

To assess the effectiveness of HDP+ in detecting PD events, we conducted experiments using real-world cryptocurrency trading data. However, building a dataset of PD events is a formidable challenge, necessitating longitudinal monitoring of the groups perpetrating the scheme. This study leverages datasets provided by La Morgia *et al.* [2020] to address this obstacle. Since the dataset provided by La Morgia *et al.* [2020] is mainly based on Binance PD events, our study also focuses on this exchange to ensure consistency and comparability. By utilizing their resources, we were able to acquire 178 Binance PD datasets used in experiments, which were downloaded using their script on GitHub[7].

The dataset introduced by La Morgia *et al.* [2020] is a collection of PD events coordinated through Telegram groups.

[6] A type of order that executes only at a specified price or better.

[7] https://github.com/SystemsLab-Sapienza/pump-and-dump-dataset/blob/master/downloader.py

These events are listed in a CSV file, where each row corresponds to a single PD occurrence and serves as input to a script that retrieves historical trading data from the Binance exchange via the Cryptocurrency Exchange Trading (CCXT) library. Notably, the La Morgia *et al.* [2020] dataset does not have raw trading records versioned, they only have the data after data engineering, thereby requiring the download of it for full analysis.

To ensure comparability with the original experiments, we compared the downloaded PD events with their available versioned data and retained only the events present in both. This step was required, as the list of PD events may have been updated after their experiments. Additionally, since Binance may delist certain assets, we provide all data used in our experiments in our GitHub repository.

Each PD event in the dataset comprises a collection of trade records spanning 12 days before to 7 days after the event. Each trade record includes information such as the symbol (cryptocurrency and its pairing coin), timestamp, trade side (buy or sell), amount, price, and volume (denominated in the pairing coin, typically Bitcoin). Given the potential pairing coin volatility, our workflow also enriches the dataset using the Live Coin Watch API[8] to obtain price and volume data in US dollars, thereby mitigating potential noise. The trade records are subsequently transformed into time series by aggregating them into predefined time chunks and computing the average price and average volume. The rush order counts within each chunk.

Another important consideration regarding the evaluation metrics is the potential presence of unidentified PD events within the dataset. The ground truth labeling in the dataset is derived from the monitoring of PD groups on Telegram, ensuring that all labeled PD events are genuine. However, it is possible that other unmonitored groups also engaged in PD activities targeting the same assets and within the dataset time frames. This possibility is supported by the observation that, given the high trading volume on Binance, the set of cryptocurrencies susceptible to price manipulation is limited, increasing the likelihood of near schemes by different groups. Consequently, the presence of undetected PD events in the dataset may negatively impact the measured precision of detection methods.

The source code for both the HDP and HDP+ methodologies, along with the necessary datasets for experimental replication, is publicly available in our GitHub repository[9]. The version corresponding to the original HDP study is preserved under the tag *HDPump*, while the HDP+ implementation is organized within a separate directory named *hdp_plus*. All code was developed in the R programming language, and we encourage the use of the RStudio IDE for reproduction.

To execute the HDP+ experiments, you may clone the repository, set the R working directory to repository root, configure the desired experimental parameters, execute the main script and invoke the main function. Additionally, we recommend registering a Live Coin Watch API key as the *LCW_API_KEY* environment variable to enable data enrichment. In the absence of the key, experiment executions will be limited to chunk sizes for which cached conversion data is already available in the repository.

The experimental setup includes four main parameters: *time frame*, *chunk size*, *preprocessing*, and *detector*. The *time frame* parameter specifies a temporal window in seconds, which constrains the length of the time series. The *chunk size* parameter defines the size of the aggregation window used to convert raw trading records into time series; within each chunk, price values are aggregated using the mean, volume using the sum, and rush orders using the count function. The *preprocessing* parameter specifies the technique to be applied to the time series data, selected from those defined in the preprocessing module. Finally, the *detector* parameter indicates the detection strategy to be used, as implemented in the strategy module.

The experiments reported in this section were conducted using a *time frame* parameter of 180,000 seconds, corresponding to a 100-hour window centered around each PD event, and a *chunk size* of 3,600 seconds (i.e., one hour). The choice of a one-hour chunk size maintains consistency with the original HDP study, which also used this configuration to enable comparison with the statistical anomaly detection approach proposed by Kamps and Kleinberg [2018]. The 100-hour time frame aligns with the experiments of La Morgia *et al.* [2020], who limited their analysis to a 24-hour window, while also satisfying the requirement of some detection algorithms for a larger number of observations. The results obtained from these experiments are presented in Table 4.

Table 4. Comparison between HD Pump Plus (HDP+), HD Pump (HDP) [Moura *et al.*, 2024] and Kamps and Kleinberg [2018] method results reported by La Morgia *et al.* [2020] for pump-and-dump event detection using 1 hour chunk size

Classifier	Precision	Recall	F1-score
Kamps (Initial)	15.6%	96.7%	26.8%
Kamps (Balanced)	38.4%	93.5%	54.4%
Kamps (Strict)	50.1%	75.0%	60.5%
HDP	56.8%	84.3%	67.9%
HDP (AD)	71.6%	46.1%	56.1%
HDP+ (AD)	90.3%	92.7%	91.5%
HDP+ (CPD)	47.7%	90.4%	62.4%
HDP+	96.4%	89.3%	92.7%

To evaluate the performance of the proposed methods, we employed the metrics of accuracy, precision, recall, and F1-score. These metrics are widely adopted in related works from Section 3, facilitating direct comparison with existing approaches. Accuracy is defined as the proportion of correctly classified observations relative to the total number of observations. Precision measures the proportion of true positive PD detections among all instances classified as PD events. Recall, also referred to as sensitivity, quantifies the proportion of correctly identified PD instances relative to the total number of actual PD events. The F1-score represents the harmonic mean of precision and recall, providing a balanced assessment of both false positives and false negatives.

Each one of these metrics is illustrated respectively in Equations 1, 2, 3 and 4. Let TP (True Positives) denote the

[8] `https://www.livecoinwatch.com/tools/api`

[9] `https://github.com/mmoura-dev/pump-and-dump`

number of observations correctly identified as part of a PD event; TN (True Negatives) represent the number of observations correctly identified as not part of a PD event; FP (False Positives) are the non-PD observations incorrectly classified as PD; and FN (False Negatives) are the PD observations that were incorrectly classified as non-PD. These metrics are widely adopted in the literature and provide a assessment of a detection model's performance, particularly in imbalanced classification problems such as PD event detection. Finally, given that the dataset comprises PD events involving different cryptocurrencies and occurring at various points in time, all metrics are computed independently for each dataset. The final reported performance scores correspond to the average values across all datasets.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{2}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{3}$$

$$\text{F1-Score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{4}$$

As the results of the AD component in Table 4 illustrate, the use of the rush order feature significantly improves the overall results achieved by the original HD Pump. The results of the CPD component showed slight improvements, which can be attributed to an update of the R package Harbinger to a newer version. The updated version 1.1.707 included several fine-tuning adjustments to its detection methods compared to version 1.0.787 used in the original HDP [Moura *et al.*, 2024]. These changes affected all results in the current work, as the updated Harbinger was applied across the entire HDP+ method. Moreover, the overall precision of HDP+ improves significantly due to the ensemble approach, which effectively combines AD and CPD to reduce false positives.

The comparison reveals a significant performance improvement. The CPD component alone achieves a 90% recall rate while maintaining a considerable advantage in precision compared to the more recall-focused configurations of the Kamps and Kleinberg [2018] study. The ensemble approach of HDP+ demonstrates a significant increase in precision compared to the standalone AD component. This result suggests that integrating CPD helps filter out market volatility-induced anomalies, improving classification confidence with only slightly reduced recall.

Despite its effectiveness, one limitation of the proposed methodology concerns the timeliness of detection. HDP+ is designed as an offline detector that relies on both pre and post-event data to identify PD schemes. However, PD schemes typically unfold over short time frames, and their market effects often dissipate within a matter of minutes [Xu and Livshits, 2019]. As a result, the current implementation of HDP+ lacks the capability for real-time detection and, therefore, cannot provide timely warnings to market participants.

6 Conclusion

This study presents HDP+, an enhanced hybrid detection method for identifying PD schemes in cryptocurrency markets. By integrating anomaly detection with change point detection techniques, HDP+ builds upon the original HD Pump method proposed by Moura *et al.* [2024] and incorporates the rush order feature introduced by La Morgia *et al.* [2020]. The proposed method significantly improves the precision of the detector compared to the original HD Pump while maintaining the advantages of statistical methods, allowing it to perform competitively against machine learning-based approaches without requiring training, even though related works are performed at different time resolutions. Additionally, we publicly release the HDP+ implementation along with the dataset used in our experiments, including the enriched version.

We aim to contribute to the detection and prevention of PD schemes that undermine the cryptocurrency market's integrity and reputation. Future research directions include refining the detection process by reducing the chunk size used for grouping trade records, thereby enhancing the granularity and timeliness of detection. Exploring additional PD detection features, such as variations of the rush order metric, could enhance accuracy. However, its applicability may be constrained by differences in data availability across exchanges.

Another promising direction is extending hybrid detection methods like HDP+ to traditional financial markets, reinforcing defenses against market manipulation beyond cryptocurrencies. Additionally, advancing online detection techniques could enable real-time identification and mitigation of PD schemes, equipping investors with timely and actionable insights. These advancements would enhance the robustness of PD detection systems, fostering more secure and transparent trading environments.

Acknowledgements

The authors thank CAPES, CNPq, and FAPERJ for partially supporting this paper.

Authors' Contributions

M.S. Moura contributed to the conception of this study. L. Baroni reviewed the text. E. Ogasawara and D.S. Mendonça advised the work and reviewed the text. M.S. Moura is the main contributor and writer of this manuscript. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence this paper.

Availability of data and materials

The datasets analyzed during the study are available at `https://github.com/mmoura-dev/pump-and-dump`.

References

Bello, A. S., Schneider, J., and Di Pietro, R. (2023). LLD: A Low Latency Detection Solution to Thwart Cryptocurrency Pump & Dumps. In *2023 IEEE International Conference on Blockchain and Cryptocurrency, ICBC 2023*. DOI: 10.1109/ICBC56567.2023.10174922.

Han, J., Pei, J., and Tong, H. (2022). *Data Mining: Concepts and Techniques*. Morgan Kaufmann, Cambridge, MA, 4th edition edition.

Hu, S., Zhang, Z., Lu, S., He, B., and Li, Z. (2023). Sequence-Based Target Coin Prediction for Cryptocurrency Pump-and-Dump. *Proc. ACM Manag. Data*, 1(1):6:1–6:19. DOI: 10.1145/3588686.

Jalal, R. N.-U.-D., Alon, I., and Paltrinieri, A. (2021). A bibliometric review of cryptocurrencies as a financial asset. *Technology Analysis and Strategic Management*. DOI: 10.1080/09537325.2021.1939001.

Kamps, J. and Kleinberg, B. (2018). To the moon: defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, 7(1). DOI: 10.1186/s40163-018-0093-5.

Kethineni, S. and Cao, Y. (2020). The Rise in Popularity of Cryptocurrency and Associated Criminal Activity. *International Criminal Justice Review*, 30(3):325 – 344. DOI: 10.1177/1057567719827051.

Kramer, D. (2005). The Way It Is and the Way It Should Be: Liability Under §10(b) of the Exchange Act and Rule 10b-5 Thereunder for Making False and Misleading Statements as Part of a Scheme to "Pump and Dump" a Stock. *University of Miami Business Law Review*, 13(2):243.

La Morgia, M., Mei, A., Sassi, F., and Stefa, J. (2020). Pump and Dumps in the Bitcoin Era: Real Time Detection of Cryptocurrency Market Manipulations. In *Proceedings - International Conference on Computer Communications and Networks, ICCCN*, volume 2020-August. DOI: 10.1109/ICCCN49398.2020.9209660.

La Morgia, M., Mei, A., Sassi, F., and Stefa, J. (2023). The Doge of Wall Street: Analysis and Detection of Pump and Dump Cryptocurrency Manipulations. *ACM Transactions on Internet Technology*, 23(1). DOI: 10.1145/3561300.

Moura, M. S., Baroni, L., Ogasawara, E., and Mendonça, D. S. (2024). HD Pump: A Hybrid Detection Approach for Pump-and-Dump Schemes in Cryptocurrency Exchanges. In *Simpósio Brasileiro de Banco de Dados (SBBD)*, pages 757–763. SBC. DOI: 10.5753/sbbd.2024.243293.

Ogasawara, E., Salles, R., Porto, F., and Pacitti, E. (2025). *Event Detection in Time Series*. Synthesis Lectures on Data Management. Springer Nature Switzerland, Cham, 1 edition. DOI: 10.1007/978-3-031-75941-3.

Rajaei, M. J. and Mahmoud, Q. H. (2023). A Survey on Pump and Dump Detection in the Cryptocurrency Market Using Machine Learning. *Future Internet*, 15(8). DOI: 10.3390/fi15080267.

Salles, R., Escobar, L., Baroni, L., Zorrilla, R., Ziviani, A., Kreischer, V., Delicato, F., Pires, P. F., Maia, L., Coutinho, R., Assis, L., and Ogasawara, E. (2020). Harbinger: Um framework para integração e análise de métodos de detecção de eventos em séries temporais. In *Anais do Simpósio Brasileiro de Banco de Dados (SBBD)*, pages 73–84. SBC. DOI: 10.5753/sbbd.2020.13626.

Schmitt, W. (2024). Bitcoin trading volumes surge after debut of long-awaited US ETFs. *Financial Times*.

Shumway, R. H. and Stoffer, D. S. (2017). *Time Series Analysis and Its Applications: With R Examples*. Springer.

Souza, J., Paixão, E., Fraga, F., Baroni, L., Alves, R. F. S., Belloze, K., Dos Santos, J., Bezerra, E., Porto, F., and Ogasawara, E. (2024). REMD: A Novel Hybrid Anomaly Detection Method Based on EMD and ARIMA. In *Proceedings of the International Joint Conference on Neural Networks*. DOI: 10.1109/IJCNN60899.2024.10651192.

Steinmetz, F., von Meduna, M., Ante, L., and Fiedler, I. (2021). Ownership, uses and perceptions of cryptocurrency: Results from a population survey. *Technological Forecasting and Social Change*, 173. DOI: 10.1016/j.techfore.2021.121073.

Takeuchi, J.-I. and Yamanishi, K. (2006). A unifying framework for detecting outliers and change points from time series. *IEEE Transactions on Knowledge and Data Engineering*, 18(4):482 – 492. DOI: 10.1109/TKDE.2006.1599387.

Victor, F. and Hagemann, T. (2019). Cryptocurrency pump and dump schemes: Quantification and detection. In *IEEE International Conference on Data Mining Workshops, ICDMW*, volume 2019-November, pages 244 – 251. DOI: 10.1109/ICDMW.2019.00045.

Wursthorn, M. (2021). A Bitcoin ETF Is Here. What Does That Mean for Investors? *Wall Street Journal*.

Xu, J. and Livshits, B. (2019). The anatomy of a cryptocurrency pump-and-dump scheme. In *Proceedings of the 28th USENIX Conference on Security Symposium*, SEC'19, pages 1609–1625, USA. USENIX Association.