CEFET/RJ
DIRETORIA DE PESQUISA E PÓS-GRADUAÇÃO

HYBRID ANOMALY AND CHANGE POINT DETECTION FOR PUMP-AND-DUMP SCHEMES IN CENTRALIZED CRYPTOCURRENCY EXCHANGES

Matheus S. Moura

Dissertação de Mestrado apresentada ao Programa de Pós-graduação em Ciência da Computação, Centro Federal de Educação Tecnológica Celso Suckow da Fonseca CEFET/RJ, como parte dos requisitos necessários à obtenção do grau de mestre.

Orientadores:
Diogo S. Mendonça

Rio de Janeiro,

April 2026

Hybrid Anomaly and Change Point Detection for Pump-and-Dump Schemes in Centralized Cryptocurrency Exchanges

Dissertação de Mestrado em Ciência da Computação, Centro Federal de Educação Tecnológica Celso Suckow da Fonseca, CEFET/RJ.

Matheus S. Moura

Aprovada por:

Presidente, Prof. Diogo Silveira Mendonça, D.Sc. (orientador)

Eduardo Soares Ogasawara, D.Sc.

Igor Machado Coelho, D.Sc.

Rio de Janeiro,

April 2026

Ficha catalográfica elaborada pela Biblioteca Central do CEFET/RJ

M929 Moura, Matheus S.
Hybrid anomaly and change point detection for pump-and-dump schemes in centralized cryptocurrency exchanges / Matheus S. Moura. — 2026.
30f. : il. color. , enc.

Dissertação (Mestrado) Centro Federal de Educação Tecnológica Celso Suckow da Fonseca, 2026.
Bibliografia : f. 28-30
Orientador: Diogo S. Mendonça

1. Criptomoedas. 2. Fraude - Prevenção. 3. Análise de séries temporais. I. Mendonça, Diogo S. (Orient.). II. Título.

CDD 332.4048

Elaborada pela bibliotecária Tania Mello – CRB/7 nº 5507/04

DEDICATÓRIA

Aos meus pais, pela paciência, compreensão e amor incondicional durante este período.

AGRADECIMENTOS

Agradeço ao meu orientador pelo constante suporte, e a todos os professores do PPCIC, pelo incentivo e apoio que tornaram possíveis as publicações ao longo desta jornada.

RESUMO

Hybrid Anomaly and Change Point Detection for Pump-and-Dump Schemes in Centralized Cryptocurrency Exchanges

Matheus S. Moura

Orientadores:
Diogo S. Mendonça

Resumo da Dissertação submetida ao Programa de Pós-graduação em Ciência da Computação do Centro Federal de Educação Tecnológica Celso Suckow da Fonseca CEFET/RJ como parte dos requisitos necessários à obtenção do grau de mestre.

O rápido crescimento dos mercados de criptomoedas intensificou as preocupações relacionadas a práticas de manipulação de mercado, e particularmente aos esquemas de pump-and-dump. A detecção desses esquemas permanece desafiadora devido à alta volatilidade das criptomoedas e a disponibilidade limitada de dados confiáveis. Trabalhos relacionado baseiam-se predominantemente em técnicas de detecção de anomalias, as quais frequentemente apresentam dificuldades com precisão e adaptabilidade. Neste trabalho, propomos dois métodos de detecção offline que exploram um arcabouço híbrido combinando detecção de anomalias e de pontos de mudança para a identificação de esquemas de pump-and-dump. O primeiro método, HD Pump, integra a detecção de anomalias de volatilidade em séries temporais de preços com a detecção de pontos de mudança aplicada ao volume negociado. O segundo método, HD Pump Plus, estende essa abordagem ao substituir a série temporal de preços por uma série temporal baseada em rush orders. A avaliação experimental em um conjunto de dados contendo 178 eventos confirmados de pump-and-dump da exchange Binance mostra que o HD Pump Plus supera abordagens estatísticas anteriores, alcançando precisão de 96,4%, recall de 89,3% e F1-score de 92,7%. Esses resultados demonstram a efetividade da estratégia híbrida de detecção em avançar o estado da arte, mantendo simplicidade metodológica.

Palavras-chave:
Criptomoedas, Pump-and-dump, Detecção de fraudes, Detecção de anomalias, Detecção de pontos de mudança

Rio de Janeiro,

April 2026

ABSTRACT

Hybrid Anomaly and Change Point Detection for Pump-and-Dump Schemes in Centralized Cryptocurrency Exchanges

Matheus S. Moura

Advisors:
Diogo S. Mendonça

Abstract of dissertation submitted to Programa de Pós-graduação em Ciência da Computação - Centro Federal de Educação Tecnológica Celso Suckow da Fonseca CEFET/RJ as partial fulfillment of the requirements for the degree of master.

The rapid growth of cryptocurrency markets has intensified concerns regarding market manipulation practices, particularly pump-and-dump schemes. Detecting such schemes remains challenging due to the high volatility of cryptocurrencies and the limited availability of reliable ground-truth data. Prior work has predominantly relied on anomaly detection techniques, which often exhibit limited precision and adaptability. In this work, we propose two offline statistical methods that explore a hybrid framework combining anomaly detection and change point detection for pump-and-dump detection. The first method, HD Pump, integrates volatility anomaly detection in price time series with change point detection applied to trading volume. The second method, HD Pump Plus, extends this approach by replacing the price time series with a rush-order-based time series. Experimental evaluation on a dataset of 178 confirmed pump-and-dump events from the Binance exchange shows that HD Pump Plus outperforms prior statistical approaches, achieving a precision of 96.4%, recall of 89.3%, and F1-score of 92.7%. These results demonstrate the effectiveness of hybrid detection strategies in advancing the state of the art while maintaining methodological simplicity.

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

Rio de Janeiro,

April 2026

Contents

I Introduction **1**

II Background **4**

III Related Work **13**

IV Method **15**

V Results **21**

VI Conclusion **26**

References **28**

List of Figures

II.1 Stages of a pump-and-dump scheme in cryptocurrency markets. 5

II.2 Anomaly detection component of the HD Pump method applied to the first-order differenced (DIFF) price time series. 8

II.3 Change point detection component of the HD Pump method applied to the cumulative sum (CUSUM) volume time series. 9

II.4 Anomaly detection component of the HD Pump Plus method applied to the rush orders time series. 10

IV.1 Workflow shared by our hybrid detection methods, HD Pump and HD Pump Plus, for pump-and-dump schemes in centralized cryptocurrency exchanges. 15

IV.2 Representation of all steps involved in a detection using the HD Pump anomaly detection component. 18

IV.3 Representation of all steps involved in a detection using the HD Pump change point detection component. 18

V.1 Probability density plot of average price in U.S. dollar comprising all 178 Binance pump-and-dump events from La Morgia et al. [2020] dataset. 22

V.2 Probability density plot of number of rush orders comprising all 178 Binance pump-and-dump events from La Morgia et al. [2020] dataset. 23

V.3 An example illustrating the difficulty of detecting a pump-and-dump event solely based on the average price of the targeted cryptocurrency. The observation corresponding to the pump-and-dump event is highlighted in red. 24

List of Tables

II.1 Comparison of the rush order feature vs. first-order differentiation of the price series using a simple outlier detection method. 8

IV.1 Sample of trading records for the Everex cryptocurrency token paired with Bitcoin on the Binance exchange. 16

IV.2 Sample of the time series extracted for the Everex cryptocurrency token, generated by the time series extraction step of the HD Pump method. 17

V.1 Comparison between HD Pump Plus, HD Pump and Kamps and Kleinberg [2018] using La Morgia et al. [2020] dataset as benchmark. 25

Lista de Acrônimos

AD	Anomaly detection.
ARIMA	Autoregressive Integrated Moving Average.
BTC	Bitcoin.
CPD	Change point detection.
CUSUM	Cumulative sum.
DIFF	First-order differencing.
EMD	Empirical Mode Decomposition.
GARCH	Generalized Autoregressive Conditional Heteroskedasticity.
GFT	General Fluctuation Test.
HDP	HD Pump.
HDP+	HD Pump Plus.
PD	Pump-and-Dump.
RED	Harbinger RED algorithm.
REMD	Robust Empirical Mode Decomposition.

Chapter I Introduction

The adoption of cryptocurrencies and blockchain technologies has been on the rise since the inception of Bitcoin (BTC) by Nakamoto [2008] [Jalal et al., 2021]. By 2025, the total market capitalization of cryptocurrencies surpassed 4 trillion dollars[1], exceeding the estimated 2025 GDP of countries such as the United Kingdom, France, and Brazil [Dept., 2025]. These facts highlight the growing significance of blockchain technologies and their increasing influence on both markets and governments.

The adoption of cryptocurrencies within traditional markets accelerated with the introduction of spot exchange-traded funds (ETFs) based on cryptocurrency assets on American exchanges in early 2024 [Schmitt, 2024]. According to CoinMarketCap[2], these instruments have already reached a total capitalization exceeding 160 billion dollars in assets under management. This rapid expansion has also drawn increasing attention from governments.

In March 2025, the US government announced the creation of the "United States Digital Asset Stockpile" as a strategic reserve [House, 2025]. Additionally, governments are investing in alternative applications of blockchain technology, such as China's incorporation of blockchain into its 14th five-year plan to strengthen national industry [Cai et al., 2021], and the Brazilian Central Bank had its own plan to develop a digital currency named Drex[3].

Despite the rapid global development and adoption, the egalitarian and decentralized nature of the blockchain ecosystem has also facilitated various forms of criminal activity. One of the earliest use cases of cryptocurrencies was their use in anonymous online black markets, such as the Silk Road drug marketplace [Hout and Bingham, 2013]. As adoption and market capitalization have expanded, illicit activities involving cryptocurrencies have also intensified. According to the Chainalysis 2025 Crypto Crime Report[4], wallets linked to illicit activities received approximately 40 billion dollars in 2024, primarily originating from scams, malware, and stolen funds. Among these illicit practices, a particularly prevalent type of scam is the Pump-and-Dump (PD) scheme.

PD is a kind of price manipulation scheme which aims to artificially inflate an asset price to subsequently sell it to unsuspecting investors, who are lured by coordinated hype created by the

[1] According to CoinMarketCap (`https://coinmarketcap.com/charts/`) and Forbes (`https://www.forbes.com/digital-assets/crypto-prices/?sh=7065bd802478`), accessed on September 15, 2025

[2] `https://coinmarketcap.com/etf/`, accessed on September 15, 2025

[3] `https://www.bcb.gov.br/estabilidadefinanceira/drex`, accessed on September 17, 2025

[4] `https://go.chainalysis.com/2025-Crypto-Crime-Report.html`, accessed on September 17, 2025

scheme insiders [Victor and Hagemann, 2019]. In cryptocurrency markets, PDs are usually orchestrated through public groups on platforms that provide encrypted communication and anonymity, such as Telegram and Discord. When the group becomes large enough to have economic power to influence prices in their selected set of target exchanges[5], they begin to operate.

Their workflow is simple: first they make an announcement of the target exchange and pairing coin[6], which is to allow the members to prepare their funds and allocate time for the scheme. Then, at the designated time, the targeted asset is revealed and the group members engage in coordinated buying and the luring of unsuspecting investors on social media [Xu and Livshits, 2019]. The pump generally reaches its peak within minutes, after which rapid sell-offs occur once signs of decline emerge, leaving unwitting holders with significant losses.

The rapid and artificial price surges induced by PD schemes diverge substantially from the usual patterns observed in cryptocurrency time series. Prior research has predominantly employed anomaly detection techniques, although these methods frequently encounter limitations in terms of precision and adaptability [Rajaei and Mahmoud, 2023]. To overcome these challenges, the present study explores hybrid anomaly and change point detection methods, assessing whether this kind of framework enhances the effectiveness of PD detection relative to approaches based exclusively on anomaly detection.

This study builds upon two previous studies of hybrid methods for PD detection, which combine anomaly detection and change point detection techniques. The first method is the published work HD Pump (HDP) that represents the initial effort to employ such a hybrid framework [Moura et al., 2024]. It is followed by HD Pump Plus (HDP+), an extension on the original method that modifies the Anomaly detection (AD) component to use a preprocessed series derived from order counts, referred to as rush orders. Both methods operate by processing raw trading records obtained from a centralized cryptocurrency exchange and applying an ensemble of offline event detection algorithms to produce a time series in which PD events are flagged.

The HDP is an offline statistical method that consists of two primary components: the AD component, which identifies volatility anomalies in the price time series, and the Change point detection (CPD) component, which analyzes the Cumulative sum (CUSUM) of the volume time series [Moura et al., 2024]. Even though HDP is the first approach to combine AD and CPD for PD detection in cryptocurrency markets, it achieves performance superior to previous statistical methods [Moura et al., 2024].

The HDP+ represents an extension of the HDP method, enhancing the original methodology by modifying the time series analyzed by the AD component. While both methods share the same overall workflow, HDP+ replaces the ensemble of algorithms applied to the price time series in HDP

[5] An organized market or platform for trading cryptocurrencies.
[6] A cryptocurrency used to trade against another cryptocurrency.

with the Robust Empirical Mode Decomposition (REMD) algorithm applied to the rush orders time series. Rush orders, introduced by La Morgia et al. [2020], are a lightweight feature defined as the count of trading records executed at the exact same millisecond. This feature captures the behavior of participants placing large buy orders that are filled by multiple sell orders simultaneously, producing multiple trade records at the same timestamp.

The experimental evaluation of HDP+ uses a dataset of 178 confirmed PD events in the Binance[7] exchange, obtained from the dataset provided by La Morgia et al. [2020]. The results, summarized in Table V.1, demonstrate a performance improvement of HDP+ over the original HDP and other methods, achieving a precision of 96.4%, recall of 89.3%, and F1-score of 92.7%. Compared to the original HDP, HDP+ is able to improve 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.

Beyond these results, our methods present additional contributions to the PD event detection landscape. The HDP introduces the first modeling of a PD event in time series as a change point, and consequently represents the first hybrid approach combining AD and CPD. In turn, HDP+ highlights the potential of rush-order-based time series for the detection of PD events and, as a statistical method like HDP, does not require training data.

The remainder of this study is organized as follows. Chapters II and III present the theoretical background and a review of related work. Chapter IV describes the methodological framework in chronological order, detailing the evolution from HDP to HDP+. Each method is discussed individually, emphasizing the specific modifications and improvements introduced in relation to its predecessor. Chapter V outlines the experimental protocol and presents the experimental results. Finally, Chapter VI concludes the paper and discusses potential future research directions.

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

Chapter II 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]. However, 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 II.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].

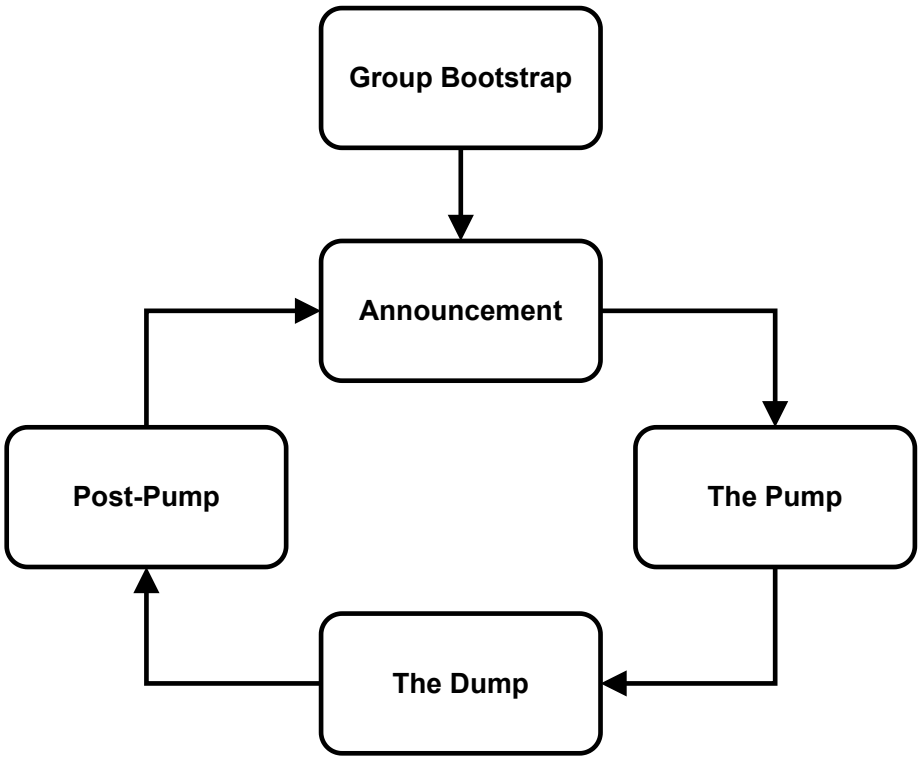


Figure II.1: Stages of a pump-and-dump scheme in cryptocurrency markets. The process begins with recruiting and organizing a pump-and-dump group and then announcing the pump details. During the pump phase, coordinated buying artificially inflates the price, often attracting unsuspecting investors. Once 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.

Classification-based anomaly detection methods have been the predominant approach to detecting PD schemes in cryptocurrency markets, which include 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 cryptocurrency. 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 commonly used methodology framework in financial markets for examining the temporal dynamics of variables such as price and trading volume [Han et al., 2022]. A

time series Z is defined as an ordered set of T observations $Z = (z_1, z_2, \ldots, z_T)$, where z_1 and z_T are the first and last observations, respectively. The observations are normally recorded at regular time intervals. The regularity of observation is related to the sampling rate of the time series. Time series with a higher sampling rate are referred to as high-frequency time series, while time series with a lower sampling rate are referred to as low-frequency time series. Time series can also be classified based on their dimensionality. A univariate time series is a time series representing one variable, while a multivariate time series is one representing multiple variables. Time series can also be classified based on their analytical components. A time series can be defined as an aggregation of three components: the trend component (β_t), the seasonal component (π_t), and the residual or random component (ω_t) [Ogasawara et al., 2025].

In the context of cryptocurrency markets, one such use for time series analysis is in detecting the sudden and artificial changes in the valuation and trading activities that can be characteristic of a PD scheme. Event detection is a specific type of time series analysis that is specifically tailored towards the detection of important changes in the behavior of a time series that can be characteristic of a specific type of event. A difference is that an event represents a type of meaningful fluctuation in the underlying process that is generating the data [Ogasawara et al., 2025]. Some other such events that can be detected in a time series include detecting network intrusions in network traffic data, equipment failures in industrial data, and diseases in electrocardiogram data.

There are three main types of events that can be detected in a time series. The first type is that of an anomaly. An anomaly represents a type of data point that is significantly different from the expected data point. The second type is that of a change point. A change point represents a type of shift between different regimes in a time series, which generate changes in the statistical properties. The third type is that of a motif. A motif represents a type of recurring data pattern that can be characterized by a specific type of shape [Ogasawara et al., 2025].

Event detection methods can be applied under different temporal settings depending on the requirements of the application. Offline detection analyzes previously collected data to identify events retrospectively. Online detection processes streaming observations in real time, often requiring incremental updates to the detection model. In contrast, event prediction seeks to forecast the occurrence of future events based on historical observations. Formally, an event may be defined as an observation at time t whose temporal characteristic deviates from the expected behavior derived from neighboring observations beyond a predefined threshold σ [Ogasawara et al., 2025].

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 [Rajaei and Mahmoud, 2023]. 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.

Event detection methods also relies on data preprocessing to improve their performance. There are many tasks that data preprocessing can be applied to time series [Ogasawara et al., 2025]. However, our hybrid methods have a wide use of data transformations. The HDP employs the First-order differencing (DIFF) and CUSUM transformations, while HDP+ uses rush orders paired with CUSUM.

The DIFF transformation is used to improve the data quality for the volatility AD component of HDP. This transformation 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 at observation t, the observations for the DIFF series d_t are 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 II.2.

Conversely, the CUSUM serves not only to enhance data quality but also is the transformation that enables the detection of PD events as change points. As depicted in Figure II.3, the CUSUM preprocessing converts each observation into the cumulative sum of itself and all preceding observations. Formally, if Z represents the original time series, the CUSUM series at time 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].

In contrast to DIFF, the rush order transformation introduced by La Morgia et al. [2020] is able to significantly improve the performance of AD for PD events as illustrated in Table II.1 [La Morgia et al., 2023]. Its concept 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

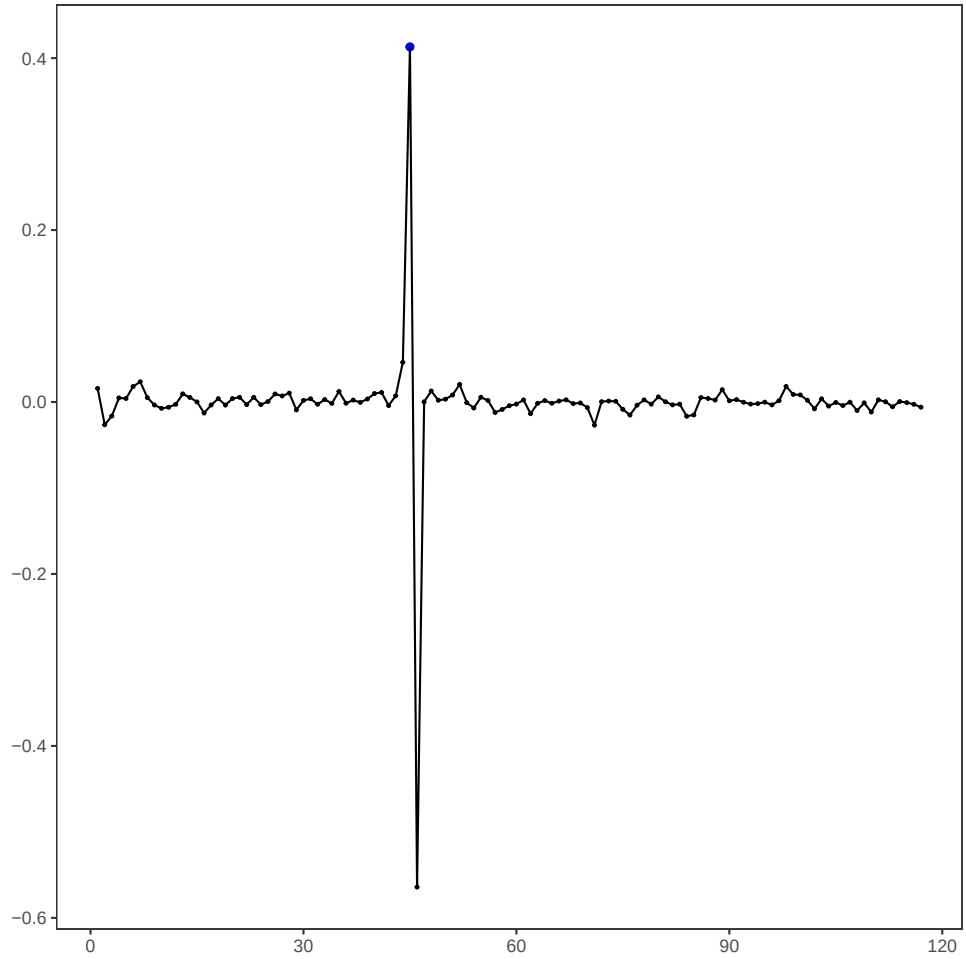


Figure II.2: Visualization of the anomaly detection component of the HD Pump method applied to the first-order differenced (DIFF) price time series of the pump-and-dump event involving the Everex cryptocurrency on September 1, 2021. The experiment uses a chunk size of 3600 seconds and a total time frame of 180000 seconds.

records. As a result, a single limit order[1] executed at multiple price levels is recorded as several individual trade records occurring at the same millisecond. The rush order transformation is thus defined as the number of times trades are executed at the exact same millisecond, as illustrated by Figure II.4.

Table II.1: 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%

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

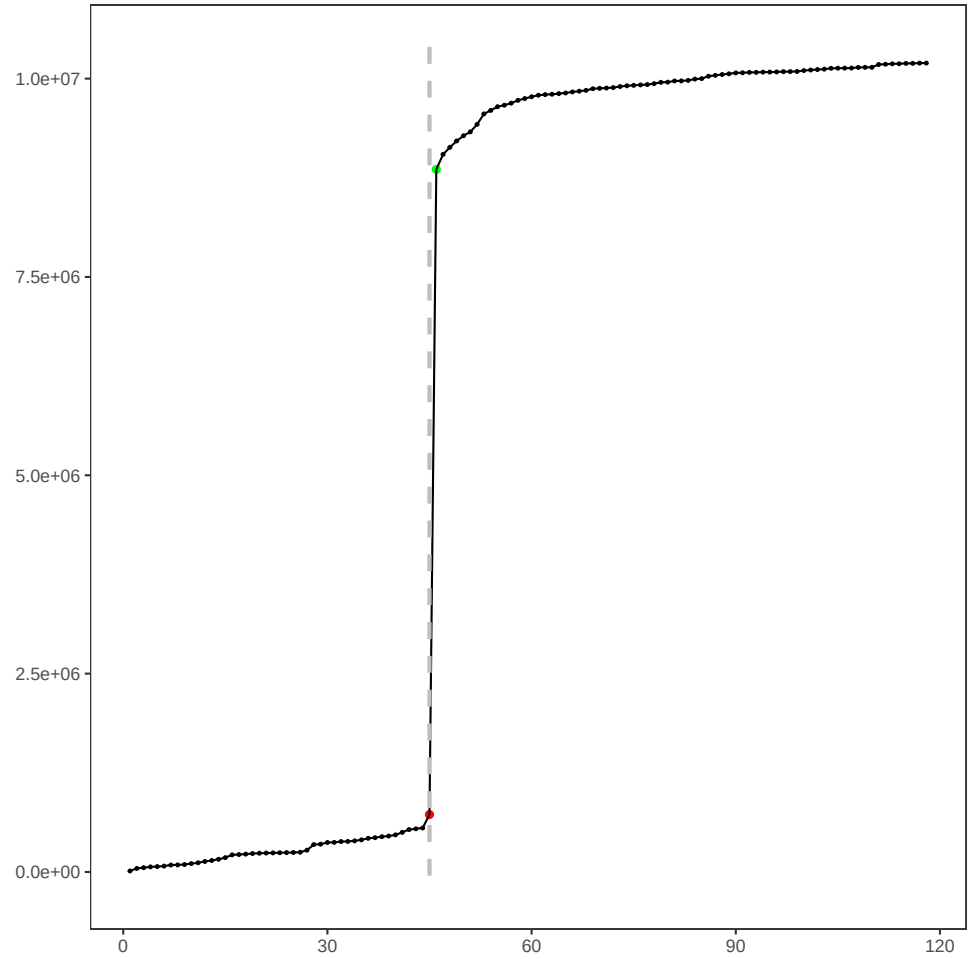


Figure II.3: Visualization of the change point detection component of the HD Pump method applied to the cumulative sum (CUSUM) volume time series of the pump-and-dump event involving the Everex cryptocurrency on September 1, 2021. The experiment uses a chunk size of 3600 seconds and a total time frame of 180000 seconds.

Our hybrid methods are also built using off-the-shelf event detection algorithms from the Harbinger framework [Salles et al., 2020]. The HDP uses an ensemble of Generalized Autoregressive Conditional Heteroskedasticity (GARCH) and Harbinger RED algorithm (RED) for anomaly detection and an ensemble of Chow test and General Fluctuation Test (GFT) for change point detection, while HDP+ swaps the ensemble of AD in favor of the REMD algorithm.

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. An observation is classified as anomalous when its actual volatility deviates significantly from the GARCH's predicted value [Ogasawara et al., 2025].

In contrast, the RED algorithm is designed as a wrapper of 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 a set of intrinsic mode functions

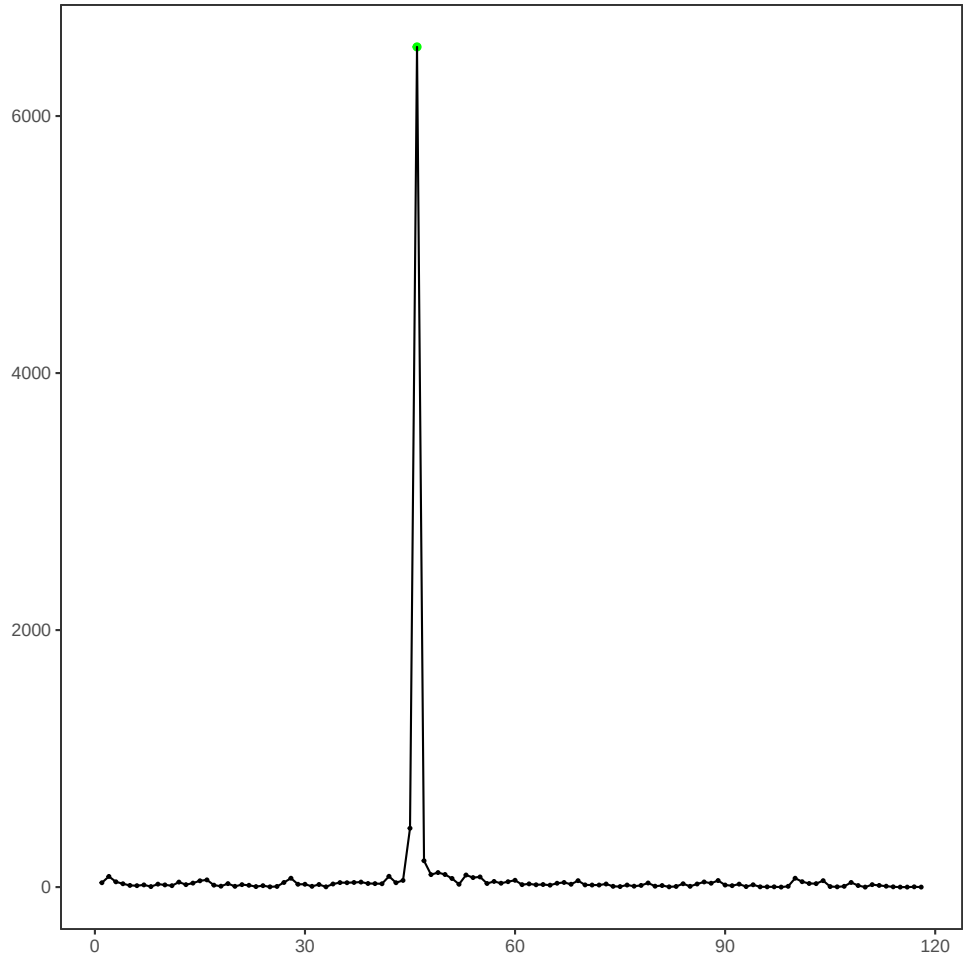


Figure II.4: Visualization of the anomaly detection component of the HD Pump Plus method applied to the rush orders time series of the pump-and-dump event involving the Everex cryptocurrency on September 1, 2021. The experiment was conducted using a chunk size of 3600 seconds and a total time frame of 180000 seconds.

(IMFs) and employs a frequency separation technique based on roughness curve analysis to distinguish between high-frequency and low-frequency components, finally detecting the anomaly on the residual component [Ogasawara et al., 2025].

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 [Ogasawara et al., 2025].

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.

On the other hand, HDP+ uses the REMD, which comprises the EMD algorithm and the Autoregressive Integrated Moving Average (ARIMA) model. The EMD is used to decompose the time series, 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, where observations classified as outliers within this probability distribution are flagged as anomalies [Souza et al., 2024].

Another important aspect to discuss is the evaluation of our hybrid methods, we employ the metrics: accuracy, precision, recall, and F1-score. These metrics are widely adopted in related works from Chapter III, 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 II.1, II.2, II.3 and II.4. Let TP (True Positives) denote the 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 an 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.

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

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

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

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

Chapter III 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. Some studies place a stronger emphasis on the theory and the problem definition, with the work of Xu and Livshits [2019] being among the most relevant. Their study provides an analysis of the behavior of PD groups by monitoring more than 300 Telegram channels associated with such schemes. However, their objective is to predict which cryptocurrency will be targeted prior to a PD group announcement, which differs from the focus of this work. Accordingly, we concentrate on studies that detect PD schemes based on time series data extracted directly from cryptocurrency exchanges.

The work of Kamps and Kleinberg [2018] is 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.

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 order feature. Chapter IV 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 previously mentioned research relied on statistical or machine learning–based anomaly detection methods, Bello et al. [2023] introduced a novel deep learning–based approach for identifying PD events in cryptocurrency markets. Their method, termed Low Latency Detection (LLD), employs an unsupervised deep learning framework designed to automatically detect PD activities on centralized exchanges. In contrast to prior studies that frequently employ supervised models dependent on future data and are thus constrained to post hoc analysis, the LLD approach emphasizes low-latency detection, prioritizing minimal detection lag to enable the timely activation of algorithmic regulatory mechanisms. Their model underwent an evaluation with 55 confirmed pump events from Binance, between January 2021 and August 2022. Across all 55 events, the method achieved a precision of 78%, a recall of 83%, and an F1-score of 80%.

However, to ensure a fair comparison, we evaluate our methods only against the method proposed by Kamps and Kleinberg [2018]. This choice is motivated by the availability of the large dataset provided by La Morgia et al. [2020] and by the fact that they have already reproduced the experiment of Kamps and Kleinberg [2018], thereby addressing a key limitation of the original study, namely the absence of a dataset with ground-truth labels. Furthermore, both of our methods operate at the same temporal resolution (referred to as the chunk size in Chapter V) using one-hour aggregation intervals. This ensures that the comparison is performed on identical time series, in contrast to the experiments of La Morgia et al. [2020], which employ finer temporal resolutions. Finally, both our methods and the approach of Kamps and Kleinberg [2018] is an offline statistical detection methods that do not require training data, eliminating the need for additional preprocessing or model training steps during evaluation.

Chapter IV Method

This chapter presents the methodology adopted to investigate the detection of PD schemes in cryptocurrency markets. The objective is to evaluate whether a hybrid framework that combines AD and CPD can improve the effectiveness of identifying PD events when compared to approaches that rely exclusively on anomaly detection. To achieve this objective, we present two offline statistical methods, namely HDP and HDP+. The first method integrates volatility-based anomaly detection applied to price time series with change point detection performed on trading volume. The second method extends this approach by replacing the price-based analysis with a rush-order-based time series, aiming to improve the detection capability of the framework. Ultimately, the proposed methodology seeks to advance the state of the art in PD detection while maintaining methodological simplicity and supporting the identification and prevention of these fraudulent practices in cryptocurrency markets.

Both HDP and HDP+ follow a similar workflow, as illustrated in Figure IV.1. 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.

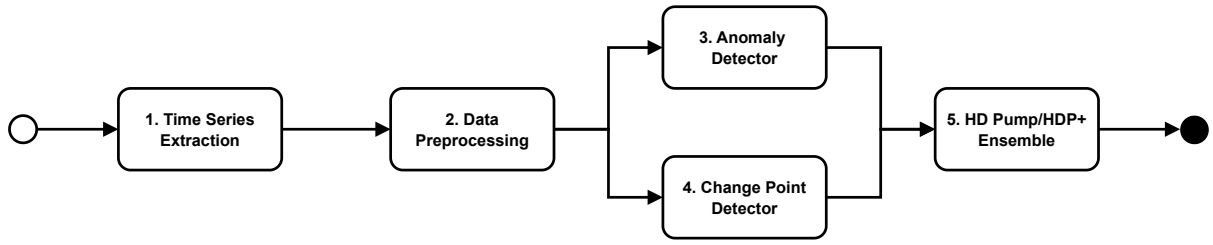


Figure IV.1: Workflow shared by our hybrid detection methods, HD Pump and HD Pump Plus, for pump-and-dump schemes in centralized cryptocurrency exchanges.

The first stage of the workflow is the transformation of input data from cryptocurrency exchanges into time series, enabling the application of event detection methods. Both methods focus on data from the Binance exchange, which, at the time of this study, is 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 correspond

to the chronological sequence of executed orders on the exchange platform within a given time interval. As illustrated in Table IV.1, each record contains several attributes describing an individual trade. The *Coin Pair* column indicates the traded market, composed of the target cryptocurrency and its pairing asset, which is not necessarily a cryptocurrency but is most commonly Bitcoin. The *Datetime* column specifies the exact timestamp at which the trade was executed. The *Side* column indicates the trade direction, identifying whether the operation corresponds to a buy or sell order. The *Price (BTC)* column represents the unit price of the traded asset denominated in the pairing asset. The *Amount* column indicates the quantity of the cryptocurrency exchanged in the trade, while the *Volume (BTC)* column represents the total traded value expressed in the pairing asset. In some datasets, the price may not be explicitly provided, as it can be derived by dividing the traded volume by the traded quantity.

Table IV.1: Sample of trading records for the Everex cryptocurrency token paired with Bitcoin on the Binance exchange.

Coin Pair	**Datetime**	**Side**	**Price (BTC)**	**Amount**	**Volume (BTC)**
EVX/BTC	2020-12-28T21:33:52.620Z	BUY	1.073e-05	3680	0.03949
EVX/BTC	2020-12-28T21:33:56.493Z	SELL	1.068e-05	52	0.00055
EVX/BTC	2020-12-28T21:33:58.183Z	SELL	1.068e-05	713	0.00761

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.

As illustrated in Table IV.2, each row corresponds to an aggregated interval. The *Bin* column identifies the starting timestamp of the time interval to which the records were aggregated. The *Number of Orders* column represents the count of trading records executed within that interval. The *Average Price (USD)* column indicates the mean trading price of the asset during the interval, expressed in US dollars. The *Volume Sum (USD)* column represents the total traded volume accumulated within the interval, also expressed in US dollars. In the HDP method, for each chunk, the average price and total traded volume are computed, resulting in the extracted time series. Additionally, the time series extraction step enriches the data by converting price and volume from Bitcoin into US dollars in order to mitigate the noise generated by the high volatility of cryptocurrency markets.

Following the extraction of time series, the second step in the workflow is data 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 for its corresponding detection component.

Table IV.2: Sample of the time series extracted for the Everex cryptocurrency token, generated by 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	0.287	10791.83
2020-12-28 23:00:00	40	0.291	9732.60
2020-12-29 00:00:00	15	0.279	3442.75

The AD component utilizes the DIFF preprocessing, which highlights changes in volatility. While the CPD component employs the CUSUM preprocessing technique. PD events are typically analyzed using anomaly detection methods [Rajaei and Mahmoud, 2023], the application of CUSUM transforms the short-term effects of PDs into structural changes in a time series, thereby enabling the use of change point detection algorithms.

Following the preprocessing stage, the resulting time series are processed by the detection components of the HDP method. The third stage is the AD component that utilizes the DIFF average price series, while the fourth stage is the CPD component which operates on the CUSUM of the volume series. Both components employ off-the-shelf event detection algorithms provided by the Harbinger framework[1] [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 these anomalies by analyzing deviations from typical series behavior [Ogasawara et al., 2025].

The HDP AD component, illustrated in Figure IV.2, is designed to focus on 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 leaving a broader detection of events to the CPD component.

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.

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

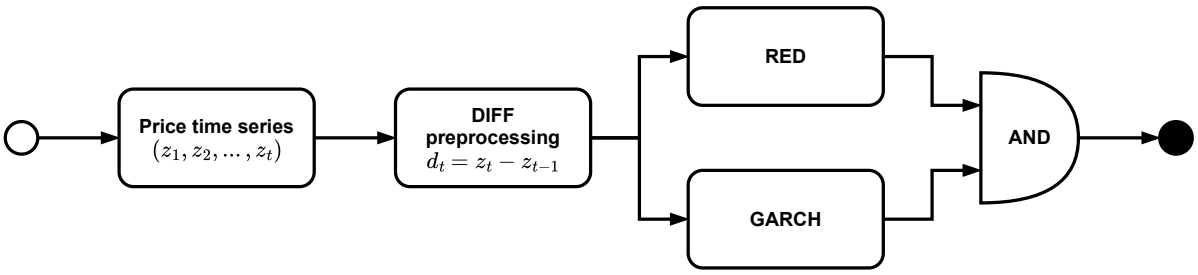


Figure IV.2: Representation of all steps involved in a detection using the HD Pump anomaly detection component.

Therefore, the CPD component employs an ensemble that integrates Chow test and GFT algorithms, as shown in Figure IV.3. The ensemble strategy is to return a positive label for observation classified as PD by both algorithms. Besides that, the Chow test algorithm used is slightly modified, for each observation labeled positive, the previous observation is also labeled as positive, emphasizing the focus of the CPD in the recall metric.

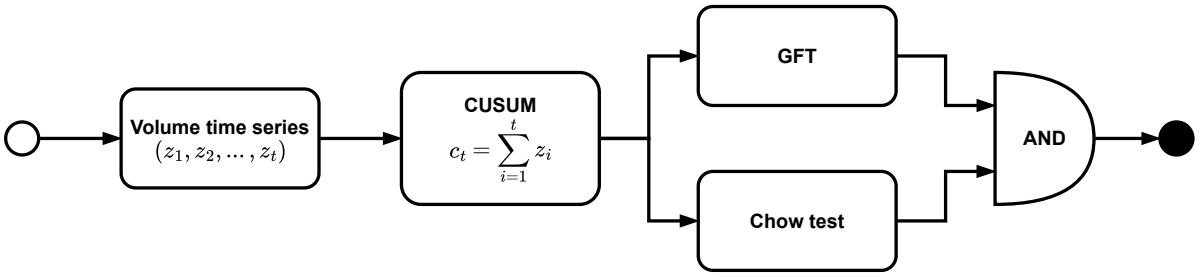


Figure IV.3: Representation of all steps involved in a detection using the HD Pump change point detection component.

The fifth and final stage of the workflow is the ensemble of the AD and CPD component results. Ensembles 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 the high confidence detections of the AD component to filter false positives, without compromising the recall achieved by the CPD component.

In order to explain the HDP method described above, a simple example is given with a hypothetical cryptocurrency token, ToyCoin, paired with USD, traded on a cryptocurrency exchange. The detection process has five steps:

The detection process starts with raw data. In the case of ToyCoin/USD pair data, the data is aggregated based on hourly intervals. The interval from 12:00 to 13:00 has 10 trades with an average price of $1.00 and a total volume of $500. The next interval from 13:00 to 14:00 has 12 trades with an average price of $1.05 and a volume of $600. A sudden event is detected from 14:00 to 15:00 with 500 trades with an average price of $5.00 and a volume of $5,000. The next interval has 400 trades with an average price of $0.50 and a volume of $4,000. The HDP method creates two time series from the above data: the average price and the total volume.

For the average price time series data, the DIFF is calculated. The DIFF is the difference between consecutive values. In the above example, the DIFF for the average price is calculated as the difference between the average prices, for example the DIFF for second interval and the third interval: $5.00 - $1.05 = $3.95. For the total volume time series data, the CUSUM is calculated. The CUSUM is the cumulative sum of the values. In the above example, the CUSUM for the total volume is calculated as the cumulative sum of the volumes: $500, $600, $5,000 and $4,000 becomes $500, $1,100, $6,100 and $10,100.

The differenced price series is then passed on for analysis by the AD component. In this case, the algorithms used are the RED and the GARCH. These algorithms detect volatility anomalies. In the toy data set provided, the high price volatility observed during the third interval is most likely to be detected by both algorithms. The CUSUM series is then passed on for analysis by a set of algorithms that belong to the CPD component. In this case, the algorithms used are the Chow test and the GFT. These algorithms detect structural changes in the statistical properties of a series. In the toy data set provided, the high rate of accumulated trading observed during the third and fourth intervals is most likely to be detected by both algorithms used in the CPD component.

The last stage involves a combination of the results from both the AD component and the CPD component. In this case, the results from both components are integrated by the HDP ensemble. This ensemble only declares a PD event if both components identify the same interval as an event. In the toy data set provided, the anomaly detection component declares the third interval anomalous, while the change point detection component declares both the third and fourth intervals as PDs. The intersection of these results is the third interval, which is then identified by the HDP as a PD event.

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 REMD algorithm. Despite these enhancements, HDP+ retains the overall workflow presented in Figure IV.1, 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.

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. The rush order feature captures this behavior by identifying bursts of trades executed at the same millisecond, indicating the simultaneous triggering of limit orders.

The rush order feature is highly effective, as highlighted in Table II.1. Its contribution to performance improvement is further reinforced by the study conducted by La Morgia et al. [2023], discussed in Chapter III, where it played a key role in achieving state-of-the-art results. The incorporation of rush orders also establishes a higher performance benchmark for the HDP+, as demonstrated by the results in Chapter V. 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 or in every circumstance. Moreover, access to more detailed order type information than Binance makes available could facilitate the development of even more precise and discriminative features, further enhancing the accuracy of PD detection.

Following that, the HDP+ introduces a new AD component in order to detect PD events in rush orders series. This new AD component relies on the REMD algorithm, from the Harbinger framework [Salles et al., 2020]. REMD is a hybrid technique that integrates EMD with an ARIMA model, which are detailed in Chapter II.

For the final classification, 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.

Chapter V Results

The experiments of HDP and HDP+ used real-world cryptocurrency trading data collected from the Binance exchange, however, building a dataset of PD events is a formidable challenge, requiring longitudinal monitoring of the groups perpetrating the scheme to have ground truth for the labels. Therefore, we leverage the dataset provided by La Morgia et al. [2020].

The La Morgia et al. [2020] dataset has 178 Binance PD events conducted by Telegram groups. The 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 using the Cryptocurrency Exchange Trading library (CCXT). However, the La Morgia et al. [2020] dataset does not have raw data versioned in their GitHub repository[1]. Besides, since March 2025 the Binance changed its API making data prior to this date available only by the portal Binance Vision, which makes their script unusable to retrieve data.

In order to facilitate the access to this dataset and the reproducibility of our work, we are providing all data used in our experiments, including La Morgia et al. [2020] raw data. Additionally, to ensure comparability with their original experiments, we filtered the events present in the dataset to retain only those used in their experiment. This step was required, as the list of PD events was updated after their experiments and Binance might remove certain cryptocurrencies from their catalog.

Each PD event in the dataset comprises a collection of trade records spanning from 12 days before to 7 days after, in one CSV file. 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). Given the potential pairing-coin volatility, our workflow also enriches the dataset using the Live Coin Watch API[2] 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.

Following the construction of the time series, we conduct an exploratory analysis of the numerical attributes *number of orders, average price (USD), sum volume (USD)* and *rush orders*. Since the

[1] `https://github.com/SystemsLab-Sapienza/pump-and-dump-dataset`
[2] `https://www.livecoinwatch.com/tools/api`

objective is to understand how these variables behave in relation to PD activity, we analyze their distributions by separating observations associated with labeled events from those corresponding to normal trading periods. We examine the density plots of each attribute, which provide a visual comparison of the distributions under both conditions.

Figure V.1 shows that the distribution of average prices during PD events resembles a half-normal distribution. This pattern arises because asset prices are strictly non-negative and because the majority of the cryptocurrencies targeted in these schemes typically trade at relatively low price levels, often only a few U.S. dollars or less. An observation derived from this distribution is that PD events do not appear to concentrate around specific target price levels. Instead, it suggests that the organizers of these schemes do not pursue a predefined price objective, but rather attempt to inflate the asset price to whatever level can be achieved under the prevailing market conditions.

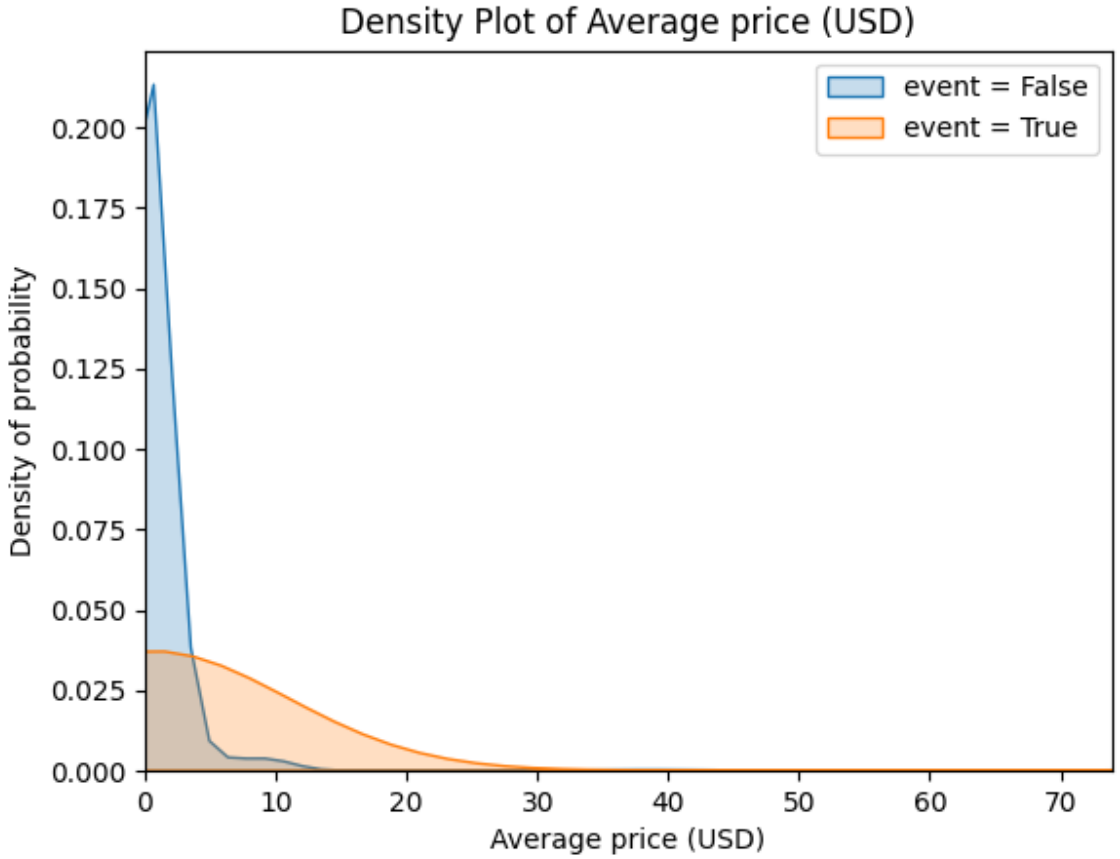


Figure V.1: Probability density plot of average price in U.S. dollar comprising all 178 Binance pump-and-dump events from La Morgia et al. [2020] dataset.

However, constructing the same probability density plots for *number of orders*, *sum volume (USD)*, and *rush orders* reveals a limitation related to the imbalance in the observations. Each dataset contains only a single labeled PD event while covering several days of regular trading activity

for the corresponding cryptocurrency. Consequently, the number of observations associated with PD periods is substantially smaller than those corresponding to normal market conditions. This imbalance results in probability density curves for the event class that are nearly flat in comparison to those of the non-event class, as illustrated in Figure V.2.

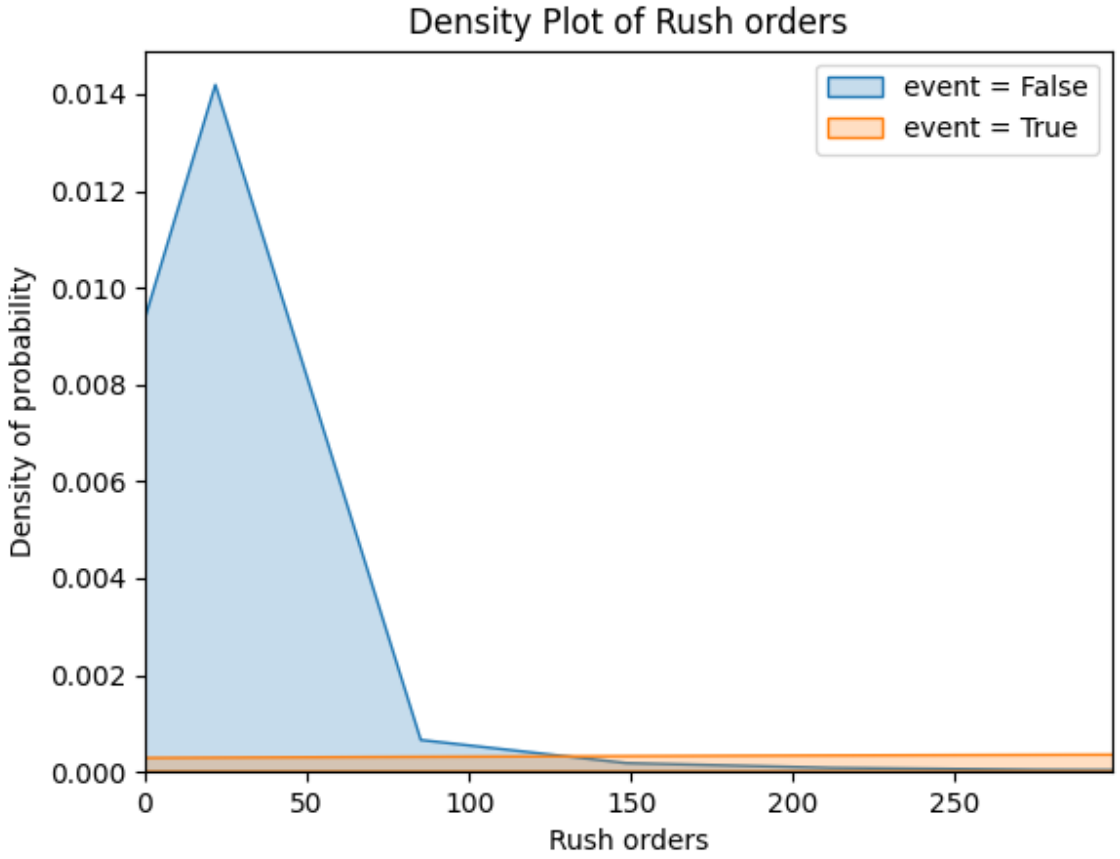


Figure V.2: Probability density plot of number of rush orders comprising all 178 Binance pump-and-dump events from La Morgia et al. [2020] dataset.

Another important consideration regarding the evaluation metrics is the potential presence of unidentified PD events within the dataset, which may be the case of Figure V.3. 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 cryptocurrencies 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, thereby increasing the likelihood of PD schemes executed by unmonitored groups. Consequently, the presence of undetected PD events in the dataset may negatively impact the measured precision.

The source code for both the HDP and HDP+ methodologies, along with the necessary datasets

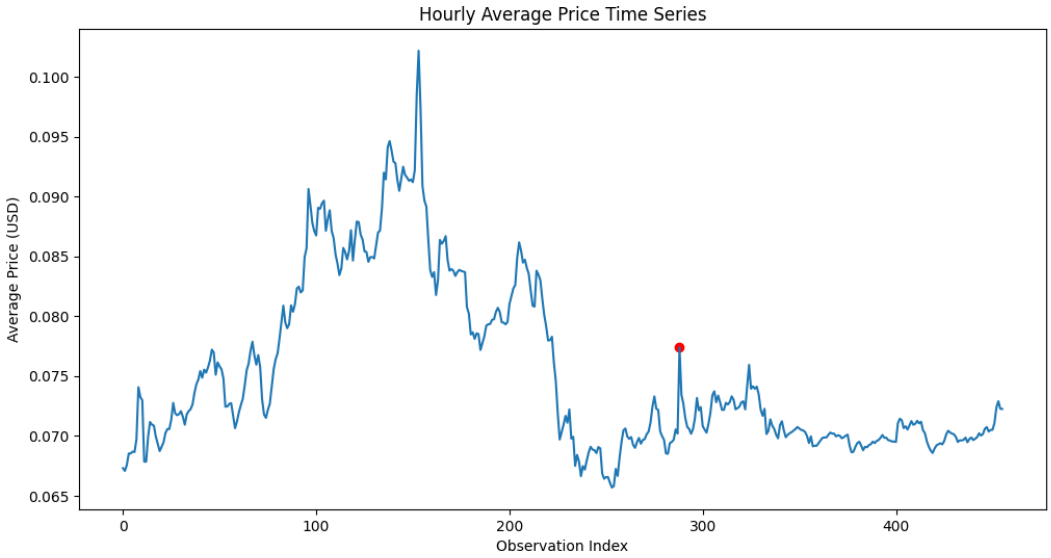


Figure V.3: An example illustrating the difficulty of detecting a pump-and-dump event solely based on the average price of the targeted cryptocurrency. The observation corresponding to the pump-and-dump event is highlighted in red.

for experimental replication, is publicly available in our GitHub repository[3]. 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*. The source code of both methods are completely in R, 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 chapter use a *time frame* parameter of 180,000 seconds, corre-

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

sponding 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 uses 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 V.1.

Table V.1: 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%

As the results of the AD component in Table V.1 illustrate, the use of the rush order feature significantly improves the overall results achieved by the original HDP. 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 1.1.787 included several fine-tuning adjustments to its detection methods compared to version 1.0.707 used in the original HDP [Moura et al., 2024]. These changes have effects on the HDP+ results, as the updated Harbinger is used in its experiments. 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.

Chapter VI Conclusion

PD schemes are a common form of market manipulation that undermine the reputation of cryptocurrency markets. Detecting such schemes remains challenging due to the high volatility of these markets. Consequently, existing state-of-the-art approaches have predominantly relied on AD techniques. In this context, this work investigates whether a hybrid method that combines AD and CPD is able to further advance the state of the art in PD detection.

To address this gap, this work presented two methods: HDP is the first iteration of a hybrid approach that combines AD and CPD for the detection of PDs in cryptocurrency markets. Subsequently, the HDP+ extends it, incorporating the analysis of rush orders. Both methods underwent peer review and publication. The HDP method was published as a short paper at the XXXIX Simpósio Brasileiro de Bancos de Dados (SBBD 2024), where it received an honorable mention in the Best Short Paper Award category [Moura et al., 2024]. The HDP+ method was accepted for publication in the Journal of Information and Data Management (JIDM), Special Issue: SBBD 2024 (Short and Extended Papers).

Both HDP and HDP+ operate in an offline detection setting, making them comparable to approaches that require access to the complete time series. Despite being a first iteration, HDP achieves performance that is comparable to other statistical offline detection methods operating on one-hour data chunks. Building upon this baseline, HDP+ yields an improvement, attaining precision and F1-score values above 90% while maintaining a recall of 89%.

These results represent an improvement over existing statistical detection methods and highlight the potential of hybrid approaches, enabling performance that is competitive with more sophisticated methods, such as those based on machine learning algorithms. Moreover, these contributions strengthen the body of knowledge on the detection and prevention of PD schemes, supporting efforts to reduce their prevalence and, consequently, to enhance the credibility of the cryptocurrency ecosystem.

However, as offline detection methods, HDP and HDP+ require access to both pre-event and post-event data to identify PD schemes, which prevents them from providing real-time warnings or timely alerts to market participants while a scheme is unfolding. Additionally, the use of one-hour data chunks for aggregation, although enabling direct comparison with earlier statistical models, also limits the temporal granularity of the detection.

Future research may address these limitations by adapting the proposed hybrid framework to an online detection setting, shifting the focus from retrospective analysis to the provision of real-time warnings for investors and exchanges. In this streaming scenario, the interaction between detection components introduces a temporal challenge: algorithms designed to detect point anomalies are likely to identify potential events earlier than change point detection algorithms. This occurs because change points represent structural changes in the statistical properties of a time series, which typically require multiple observations to confirm. Consequently, an integrated ensemble of detectors must either account for this inherent detection delay or manage the continuous revision of observations that may have already been preliminarily flagged as PD events.

Another promising direction is to extend the applicability of the approach beyond cryptocurrency markets, investigating its effectiveness in traditional financial settings, such as stock markets, including small-cap equities. Overall, this work shows that hybrid detection strategies offer a promising direction for PD detection, balancing effectiveness and methodological simplicity. The proposed methods provide a foundation upon which future online and cross-market extensions can be built, contributing to more reliable and transparent market surveillance mechanisms.

References

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

Cai, L., Sun, Y., Zheng, Z., Xiao, J., and Qiu, W. Blockchain in china. *Commun. ACM*, 64(11):88–93, 2021.

Dept., I. M. F. R. *World Economic Outlook, April 2025: A Critical Juncture amid Policy Shifts*. International Monetary Fund, USA, 2025.

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

House, T. W. Establishment of the strategic bitcoin reserve and united states digital asset stockpile. https://web.archive.org/web/20260113082517/https://www.whitehouse.gov/presidential-actions/2025/03/establishment-of-the-strategic-bitcoin-reserve-and-united-states-digital-asset-stockpile/, 2025.

Hout, M. C. V. and Bingham, T. 'silk road', the virtual drug marketplace: A single case study of user experiences. *International Journal of Drug Policy*, 24(5):385–391, 2013.

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

Jalal, R. N.-U.-D., Alon, I., and Paltrinieri, A. A bibliometric review of cryptocurrencies as a financial asset. *Technology Analysis and Strategic Management*, 2021.

Kamps, J. and Kleinberg, B. To the moon: defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, 7(1), 2018.

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

Kramer, D. 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, 2005.

La Morgia, M., Mei, A., Sassi, F., and Stefa, J. 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, 2020.

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

Moura, M. S., Baroni, L., Ogasawara, E., and Mendonça, D. S. 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, 2024.

Nakamoto, S. Bitcoin whitepaper. *URL: https://bitcoin. org/bitcoin. pdf-(: 17.07. 2019)*, 9:15, 2008.

Ogasawara, E., Salles, R., Porto, F., and Pacitti, E. *Event Detection in Time Series*. Synthesis Lectures on Data Management. Springer Nature Switzerland, Cham, 1 edition, 2025.

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

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. 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, 2020.

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

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

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

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

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

Xu, J. and Livshits, B. 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, 2019.