# HD Pump: A Hybrid Detection Approach for Pump-and-Dump Schemes in Cryptocurrency Exchanges

**Matheus S. Moura, Laís Baroni, Eduardo Ogasawara, Diogo S. Mendonça**

[1] Federal Center for Technological Education of Rio de Janeiro - CEFET/RJ

{matheus.moura,lais.baroni}@aluno.cefet-rj.br

eogasawara@ieee.org, diogo.mendonca@cefet-rj.br

***Abstract.*** *The adoption of cryptocurrencies has created a favorable environment for price manipulation practices, such as pump-and-dump (PD) schemes. These schemes aim to artificially inflate an asset's price, followed by a rapid sell-off, which may harm unaware investors. Given the brief duration of PD scheme effects, their impact on the asset's price series can be considered anomalies. Most studies rely on classification-based anomaly detection techniques to identify the PD event, which presents an opportunity to explore techniques beyond anomaly detection. To address this, we explore the combination of anomaly and change point detection to enhance pump-and-dump scheme detection. We introduce HD Pump, a hybrid detection method that integrates both techniques. Experimental results demonstrate that our hybrid approach significantly improves performance, achieving a 6.7% increase in precision and a 9.3% increase in recall compared to the benchmark method that solely uses anomaly detection.*

## 1. Introduction

The adoption of cryptocurrencies has experienced significant growth over the past decade, leading to a substantial market expansion within traditional financial sectors and in public perception [Jalal et al., 2021]. This trend is exemplified by notable milestones, 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 trading commencement of spot ETFs on American stock exchanges in early 2024 [Schmitt, 2024]. The surge in public interest can be associated with the exponential growth observed in the cryptocurrency market during late 2017 and early 2018, which captured widespread attention from mainstream media, attention that recurs with each subsequent occurrence of market rallies [Steinmetz et al., 2021]. Consequently, comprehensive studies are needed to assess the risks associated with this fairly new asset class.

The vast majority of cryptocurrencies operate without reliance on any central authority and offer a degree of anonymity. These features facilitate their use in illicit activities, including but not limited to money laundering, drug trafficking, and hacking [Kethineni and Cao, 2020]. Among the various fraudulent practices associated with cryptocurrencies, one prominent scheme is the pump-and-dump (PD) scheme, characterized by the artificial inflation of an asset's price followed by the sale of the acquired assets at a higher price [Victor and Hagemann, 2019].

In cryptocurrency markets, PD schemes are typically organized within publicly accessible groups on platforms offering encryption and anonymity, such as Telegram or

Discord. These groups recruit members until acquiring sufficient economic power to influence an asset's price. The scheme operation begins with an announcement a few days prior, detailing the exchange[1], the pairing coin[2], and the precise timing of the pump [La Morgia et al., 2020]. The targeted asset is announced at a predetermined time, prompting members to artificially buy, hold, and promote it to inflate its price artificially. The pump typically peaks within minutes, followed by rapid panic-selling at the first sign of decline, returning the cryptocurrency to its pre-pump price and trading volume within approximately half an hour [Xu and Livshits, 2019].

The brief moment of price inflation and panic-selling caused by PD schemes does not comply with the general behavior of the price and volume time series. These time series can be defined as the ordered list of observations $(z_1, z_2, \ldots, z_t)$, where $z_t$ $(t = 1, \ldots, T)$ is the observation at a particular position and $T$ is the length of the time series [Han et al., 2022]. Most existing studies in the field leverage anomaly detection methods to address the problem, particularly those based on classification using machine learning models, such as random forest and support vector machine [Rajaei and Mahmoud, 2023]. This scenario presents an opportunity to explore techniques beyond anomaly detection, such as change point detection and hybrid methods, to enhance the efficacy of pump-and-dump detection. Therefore, the research problem seeks to determine if and how the combination of anomaly and change point detection can aid PD detection.

This paper presents the Hybrid Detection of the Pump-and-Dump technique (HD Pump) for PD schemes in cryptocurrency exchanges. HD Pump integrates techniques from anomaly detection [Olteanu et al., 2023] and change point detection [Truong et al., 2020]. Besides that, it demonstrates a significant performance improvement, achieving an 84.3% recall, 56.8% precision, and 67.9% F1-score compared to the strict configuration of Kamps and Kleinberg [2018], which scored 75.0% recall, 50.1% precision, and 60.5% F1-score on the same datasets.

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 HD Pump 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 pump-and-dump 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, typically of small size and with low trading volumes, as their target. Subsequently, they accumulate significant quantities of this security. Following the acquisition, they taunt the security, disseminating false and misleading 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

[1] An organized market or center for trading cryptocurrencies.

[2] Cryptocurrency used to trade against another cryptocurrency.

[Kethineni and Cao, 2020]. These characteristics have created a favorable environment for online criminal activities using cryptocurrencies, such as PD schemes.

The standard procedure for a PD scheme on cryptocurrency exchanges involves perpetrators forming a public group and recruiting participants until a critical mass is reached. They then announce the pump details a few days prior, specifying the exchange, the pairing coin, and the event's timing [La Morgia et al., 2020]. The targeted coin is announced at a predetermined time, and participants are urged to buy and hold it, causing a rapid price increase. Shortly after, participants panic-sell, quickly returning the cryptocurrency's price and trading volume to pre-pump levels. The perpetrators then release a biased review, highlighting supposed high profitability for participants [Xu and Livshits, 2019].

## 3. Related Work

Kamps and Kleinberg [2018] was one of the pioneering studies in cryptocurrency PD schemes. It offered an early formalization of these schemes and presented a methodology for detecting them using AD algorithms. Their methodology detects local conditional point anomalies that co-occurrence price and volume anomalies. Furthermore, their method has parameters that optimize recall, precision, or a balance between metrics.

A significant challenge Kamps and Kleinberg [2018] encountered was the limited availability of labeled data at the time. This challenge was eased by the contributions of La Morgia et al. [2020], which introduced a dataset of confirmed PD schemes in the literature. Their work focused on monitoring PD groups that primarily operated on the Binance exchange[3] 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 the table 1.

Beyond the studies discussed here, there are other works in the literature that focus on the detection of PD schemes, particularly through machine learning-based approaches [Rajaei and Mahmoud, 2023]. However, we use the work of Kamps and Kleinberg [2018] as a benchmark for two primary reasons. First, both Kamps and Kleinberg [2018] and HD Pump are evaluated using the same dataset provided by La Morgia et al. [2020]. Second, both methodologies utilize the same time resolution (chunk size) of one hour. These factors ensure a fair and consistent basis for comparison between the results. Additionally, since Kamps and Kleinberg [2018] relies solely on anomaly detection, this comparison highlights the difference in employing a hybrid method such as HD Pump.

## 4. HD Pump

The HD Pump workflow for detecting PD events in cryptocurrency exchanges is outlined in Figure 1. It illustrates the stages of converting the raw data into time series, data preprocessing, anomaly detection, change point detection, and how the HD Pump generates the aggregated result. These stages define the process from raw data handling to integrating off-the-shelf technique results into the final HD Pump output.

The first step in our workflow involves transforming raw data into a time series. Cryptocurrency exchanges typically provide historical data in the form of a list of trade

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

records with a timestamp, volume, price, and whether it is a buy or sell order. These records are grouped temporally to enable time series processing, forming a time series with regular intervals. Data points are generated by calculating the average price and volume for each period, referred to as the chunk size.

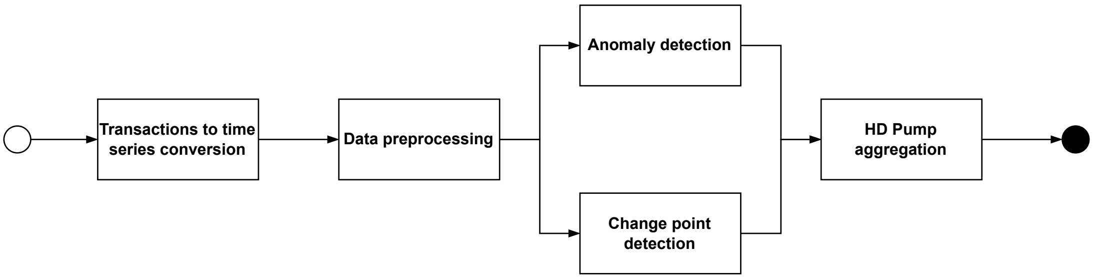

**Figure 1. HD Pump workflow for the detection of pump-and-dump**

After constructing the time series, it is essential to transform it to facilitate the application of event detection methods. The differencing preprocessing (DIFF) technique is applied to the price series to highlight changes in volatility. Formally, if $z_t$ represents the original time series at time $t$, the first-order differentiated series $d_t$ is defined as $d_t = z_t - z_{t-1}$ [Shumway and Stoffer, 2017]. This transformation accentuates abnormal fluctuations in the volatility of the price time series induced by PD schemes.

Another preprocessing technique utilized is the cumulative sum (CUSUM) applied to the volume series, which enables the observation of the effects of a PD scheme as change points rather than anomalies. 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]. These significant changes can indicate PD schemes, making CUSUM a required preprocessing step for change point detection in the context of PD event identification.

Subsequently, the transformed time series are processed by the main components of the HD Pump. The anomaly detection (AD) component identifies volatility anomalies within the price time series. In contrast, the change point detection (CPD) component is dedicated to change points in the cumulative volume series. Both components are ensembles of event detection techniques from the Harbinger framework [Ogasawara et al., 2023]. However, the specific techniques are not detailed in this work due to space constraints.

The AD component comprises an ensemble of the GARCH and RED techniques. The GARCH model is initially employed to detect volatility anomalies in the first-order differentiated price series, and the RED model is applied to the same series to detect multi-level volatility changes. The ensemble result for the AD component is derived by combining the results from both techniques; if both techniques label an observation as positive, the final result is considered positive.

The CPD component comprises an ensemble of the Chow test and the GFT applied to the CUSUM of the volume series. Initially, the Chow test is applied to the series, and its detection is modified; any observation following a positive label by the Chow test is also labeled positive. After that, the GFT model is applied to the same series. The aggregated result of the CPD component is also determined by whether the modified Chow test and GFT label an observation as positive; the final result is positive.

Then, the results of the AD and CPD components are combined using an ensemble approach. Ensemble is a robust and widely used approach for combining methods. The general procedure begins with an input dataset $X$, containing $n$ data tuples in a $d$-dimensional space. Ensemble methods first create $K$ base detectors, each operating in a random subspace $X_i$ of the original dataset. The base detectors are applied within each subspace to assign scores to the data tuples. These scores are aggregated to produce an overall score for each data tuple [Han et al., 2022].

Since the AD component is applied to a first-order differentiated series, its results should be shifted one observation to the right because the differentiation preprocessing removes the first observation. After accounting for this difference, the results can be combined. The AD component is designed for high precision, though possibly at the expense of recall. Therefore, the results ensemble logic uses its results to achieve perfect precision in easily detecting events. At the same time, the recall-focused CPD component handles any other scenario alone. Suppose the AD component labels only one observation as positive, and the CPD component also labels the same observation as positive. In that case, this is the only positive label in the aggregated result. Otherwise, the result of the CPD component is used as the aggregated result. The HD Pump method is made available at GitHub[4].

## 5. Results

Building a dataset of PD events is a formidable challenge, necessitating longitudinal monitoring of the groups perpetrating the scheme. To address this obstacle, this study leverages datasets provided by La Morgia et al. [2020]. Consequently, our investigation also focuses on PD occurrences within the Binance exchange. By utilizing their resources, we were able to acquire 178 Binance PD datasets used in experiments, which were downloaded using their script[5]. This preparatory step is required to be performed only once, as the downloaded datasets are utilized for all subsequent analyses.

Each PD dataset consists of a list of trade records from 12 days before to 7 days after the event. Each trade record includes the symbol (coin/pairing coin), timestamp, side (buy or sell), amount, price, and volume (denoted in the pairing coin, which usually is bitcoin). Given bitcoin's volatility, we enrich these records using the Live Coin Watch API[6] to obtain price and volume data in US dollars, thereby reducing potential noise. Subsequently, the trade records are transformed into time series by grouping them into chunks and calculating their average price and volume.

This process generates the time series data for both price and volume. To maintain consistency with the results reported by La Morgia et al. [2020], the datasets used in our analysis were reduced to 100 hours surrounding the event, and the chunk size for grouping

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

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

[6] https://www.livecoinwatch.com/tools/api

trade records was fixed at one hour. The results from the described analysis using the 178 PD datasets are summarized in Table 1.

**Table 1. Comparison between HD Pump 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% |
| HD Pump - AD component | 71.6% | 46.1% | 56.1% |
| HD Pump - CPD component | 47.1% | 90.0% | 61.8% |
| HD Pump (Hybrid) | 56.8% | 84.3% | 67.9% |

The HD Pump detector comprises two primary components, each with distinct objectives that can be evaluated independently. The AD component prioritizes precision over recall, effectively identifying straightforward cases. Conversely, the CPD component emphasizes recall without entirely sacrificing precision, making it suitable for more general applications. When integrated within the HD Pump method, the AD component enhances the precision of the CPD component by accurately handling easy cases, thus preserving the CPD's recall. This synergy maximizes the overall F1 score of the HD Pump method.

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—initial and balanced—of the Kamps and Kleinberg [2018] study. The HD Pump method is similar to the strict configuration, as both techniques aim to maximize the F1 score. HD Pump outperforms, scoring 84.3% recall, 56.8% precision, and 67.9% F1-score, compared to their 75.0% recall, 50.1% precision, and 60.5% F1-score. These results highlight the benefits of using a hybrid method over a plain AD method.

## 6. Conclusion

This paper presents HD Pump, a hybrid method for detecting PD schemes in cryptocurrency markets. Our findings underscore the improvements from combining anomaly detection and change point detection techniques. HD Pump is a step forward to help market actors detect and prevent market abuse and criminal behavior.

Future works could explore reducing the chunk size for grouping trade records, which may improve the granularity and timeliness of the detection process. Another direction involves applying hybrid methods like HD Pump to traditional financial markets, strengthening defenses against PD schemes beyond the realm of cryptocurrencies. Furthermore, future research can explore online detection methods, enabling real-time identification and mitigation of PD schemes and offering more immediate information for unsuspecting investors. These advancements could significantly strengthen the robustness and responsiveness of PD scheme detection systems.

## References

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

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*.

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

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

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.

Ogasawara, E., Salles, R., Lima, J., Baroni, L., Castro, A., Carvalho, L., Borges, H., Carvalho, D., Coutinho, R., Bezerra, E., Pacitti, E., and Porto, F. (2023). harbinger: A Unified Time Series Event Detection Framework.

Olteanu, M., Rossi, F., and Yger, F. (2023). Meta-survey on outlier and anomaly detection. *Neurocomputing*, 555.

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).

Schmitt, W. (2024). Bitcoin trading volumes surge after debut of long-awaited US ETFs. Technical report, https://www.ft.com/content/f30ece62-0f1c-492a-8ccd-63ec9730573c.

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

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*, 173:121073.

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.

Truong, C., Oudre, L., and Vayatis, N. (2020). Selective review of offline change point detection methods. *Signal Processing*, 167.

Victor, F. and Hagemann, T. (2019). Cryptocurrency Pump and Dump Schemes: Quantification and Detection. In *2019 International Conference on Data Mining Workshops (ICDMW)*, pages 244–251.

Wursthorn, M. (2021). A Bitcoin ETF Is Here. What Does That Mean for Investors? Technical report, https://www.wsj.com/articles/a-bitcoin-etf-is-almost-here-what-does-that-mean-for-investors-11634376601.

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